Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Exact Simulation of the SABR Model

http://orcid.org/0000-0002-3235-3340Ning Cai, Yingda Song, Nan Chen

To cite this article:
http://orcid.org/0000-0002-3235-3340Ning Cai, Yingda Song, Nan Chen (2017) Exact Simulation of the SABR Model. Operations Research 65(4):931-951. https://doi.org/10.1287/opre.2017.1617

http://pubsonline.informs.org/journal/opre/

OPERATIONS RESEARCH
Vol. 65, No. 4, July–August 2017, pp. 931–951
ISSN 0030-364X (print), ISSN 1526-5463 (online)

CONTEXTUAL AREAS

Exact Simulation of the SABR Model

Ning Cai,[a] **Yingda Song,**[b] **Nan Chen**[c]

[a] Department of Industrial Engineering and Logistics Management, The Hong Kong University of Science and Technology, Kowloon, Hong Kong; [b] Antai College of Economics and Management, Shanghai Jiao Tong University, Shanghai, China 200030; [c] Department of Systems Engineering and Engineering Management, The Chinese University of Hong Kong, Shatin, Hong Kong

Contact: ningcai@ust.hk, http://orcid.org/0000-0002-3235-3340 (NiC); songyd@sjtu.edu.cn (YS); nchen@se.cuhk.edu.hk (NaC)

Received: July 16, 2011
Revised: April 13, 2013; December 17, 2014
Accepted: January 25, 2017
Published Online in Articles in Advance: June 30, 2017

Subject Classifications: finance: asset pricing; simulation: efficiency, application
Area of Review: Financial Engineering

https://doi.org/10.1287/opre.2017.1617

Abstract. The stochastic alpha-beta-rho (SABR) model becomes popular in the financial industry because it is capable of providing good fits to various types of implied volatility curves observed in the marketplace. However, no analytical solution to the SABR model exists that can be simulated directly. This paper explores the possibility of exact simulation for the SABR model. Our contribution is threefold. (i) We propose an *exact simulation* method for the forward price and its volatility in two special but practically interesting cases, i.e., when the elasticity $\beta = 1$, or when $\beta < 1$ and the price and volatility processes are instantaneously uncorrelated. Primary difficulties involved are how to simulate two random variables whose distributions can be expressed in terms of the Hartman-Watson and the noncentral chi-squared distribution functions, respectively. Two novel simulation schemes are proposed to achieve numerical accuracy, efficiency, and stability. One stems from numerical Laplace inversion and Asian option literature, and the other is based on recent developments in evaluating the noncentral chi-squared distribution functions in a robust way. Numerical examples demonstrate that our method is fast and accurate under various market environments. (ii) When $\beta < 1$ but the price and volatility processes are correlated, our simulation method becomes a *semi-exact* one. Numerical results suggest that it is still quite accurate when the time horizon is not long, e.g., no greater than one year. For long time horizons, a *piecewise semi-exact* simulation scheme is developed that reduces the biases substantially. (iii) For European option pricing under the SABR model, we propose *a conditional simulation* method, which reduces the variance of the plain simulation significantly, e.g., by more than 99%.

Funding: The research of the first author is partially supported by the General Research Fund of the Hong Kong Research Grants Council (Project No. 16206415) and University Grants Committee grant (Project DAG12EG07-11). The research of the second author is supported by Natural Science Foundation of China [Grant 71571129]. Part of the research of the third author is supported by Hong Kong Research Grant Council via the General Research Fund Scheme (Project No. 14201114).
Supplemental Material: The e-companion is available at https://doi.org/10.1287/opre.2017.1617.

Keywords: exact simulation • SABR model • semi-exact simulation • piecewise semi-exact simulation • time change • Asian options • noncentral chi-squared distributions • Hartman-Watson distributions

1. Introduction

It has long been observed in the derivatives markets that the implied volatility is not constant across options contingent on the same underlying but with different strike prices. The implied volatility is often relatively lower for near-the-money foreign currency options, and tends to rise in the in-the-money and the out-of-the-money directions. Traders refer to this stylized pattern as the *volatility smile*. In the equity and interest rate markets, a typical shape of the implied volatility known as the *volatility skew* is that it decreases as the strike price increases. The ubiquitous existence of volatility smiles and skews poses a great challenge to the practice of risk management in fixed income and foreign exchange trading desks. These desks usually have large exposure across a wide range of strikes. The fact that different options correspond to different implied volatilities indicates that we have to model these smiles accurately to achieve an accurate and stable hedging.

Stimulated by practitioners' need to deal with the smile risk, there is a growing interest in stochastic volatility models in the area of financial engineering; see, for example, Hull and White (1987, 1988), Scott (1987), Wiggins (1987), Johnson and Shanno (1987), Stein and Stein (1991), Heston (1993), Duffie et al. (2000), Fouque et al. (2000), Barndorff-Nielsen and Shephard (2001), and so forth. The stochastic alpha-beta-rho (SABR) model introduced by Hagan et al. (2002) has gained popularity in the financial industry to model foreign exchange and interest rate markets. On the one hand, the SABR model can fit implied volatility smiles very well with only a few parameters; on the other hand, it is capable of generating correct

comovements between the underlying and the smile curves, which overcomes a salient drawback of conventional local volatility models proposed by Dupire (1994, 1997) and Derman and Kani (1994, 1998).

Despite the above attractive features of the SABR model, its complicated structure prevents us from obtaining closed-form solutions to related pricing problems of financial derivatives, including European options, swaptions, and other path-dependent options such as forward starting options. The literature so far mainly relies on partial differential equation (PDE) based approaches to find asymptotic expansions for option prices under the SABR model. Hagan et al. (2002) apply the singular perturbation technique in the PDE to obtain an explicit asymptotic expansion for the implied volatility, and then substitute it into Black's formula to yield an analytical approximation to the European option price. Obłój (2008) points out an inconsistency between the results in Hagan et al. (2002) and Berestycki et al. (2004) under a more general framework and obtains a corrected version of the implied volatility formula. Wu (2012) constructs a series expansion for the joint density of the price and the volatility for the SABR model through a hierarchy of parabolic equations and a near-Gaussian coordinate transformation.

However, the aforementioned approaches may have three limitations. First, to the best knowledge of the authors, almost all these methods are focused on computations related to the distributions of the price and the volatility, European option pricing, or implied volatilities. It might be difficult to extend these approaches to price actively traded path-dependent options such as the forward starting options. Second, the approximations obtained via these approaches perform well only under some limiting conditions of the time to maturity and strike price, and may otherwise result in relatively large deviations from the true values. For instance, the approximate implied volatilities of Hagan et al. (2002) become less accurate when the time to maturity is not small, or when the option is deep in or out of the money; see, e.g., Hagan et al. (2002), Benaim and Friz (2009) for detailed discussions. Third, in general, only the first several terms are given in the existing asymptotic expansion formulas, and the higher-order terms are always too complicated to obtain easily. Thus, it becomes difficult to control the errors in various parameter settings.

In practice, Monte Carlo simulation is a popular method for option pricing under the SABR model. Nonetheless, the conventional discretization methods for the SDE simulation (see, e.g., Chapter 6 of Glasserman 2004) may not be efficient enough to generate sufficiently accurate estimates. As shown in the numerical examples in this paper, if choosing the Euler scheme, we always have to simulate a large number of sample paths to reduce the statistical error and, at the same time, to increase the number of time steps in each path to reduce the discretization bias. This may incur a huge computational burden. More seriously, the discretization methods potentially lead to an extra distortion of the price distribution for the SABR model besides the conventional one caused by the discretization. Under the SABR model, the price process is governed by a constant elasticity of variance (CEV)-type diffusion whose value ranges in the interval $[0, +\infty)$. However, the discretization methods may generate negative values in the intermediate steps with a significant probability. Therefore we have to restrict the values within the domain by, e.g., truncating the negative values to zero. This would then bring a new distortion to the original price distribution.

In light of the absence of closed-form option pricing formulas and the drawbacks of PDE-based asymptotic expansions and conventional discretization simulation, the aim of this paper is to explore the possibility of *exact simulation* for the SABR model. Apparently, this relies heavily on the knowledge about the *exact joint distribution* of the forward price and the volatility. Unfortunately, in general cases, it seems very difficult to derive it, if not impossible. To the authors' best knowledge, there exists no "complete" answer to this question in the literature. A "partial" solution is given as follows. When the elasticity $\beta = 1$, we can show that the forward price is conditionally normal given the terminal volatility and integrated variance. For any $\beta \in [0, 1)$, if the price and volatility processes are instantaneously uncorrelated ($\rho = 0$), Islah (2009) explicitly derives the conditional distribution of the forward price, given the terminal volatility and integrated variance, when 0 is either an absorbing or a reflecting boundary; if $\rho \neq 0$, the previously exact result for the conditional distribution serves as an interesting approximation. See Propositions 2.1 and 6.1.

Based on these results, our paper proposes an exact simulation method for the forward price and the volatility of the SABR model, when either $\beta = 1$, or $\beta < 1$ and $\rho = 0$ with an absorbing or a reflecting boundary. Note that these two special cases are interesting and important for the SABR model in practice. In their original paper of the SABR model, (Hagan et al. 2002, p. 91) point out that "*The exponent β and correlation ρ affect the volatility smile in similar ways. . . . , it is difficult to distinguish between the two parameters*." Therefore, practically we may either first fix ρ (e.g., $= 0$) and then allow β to change, or first fix β (e.g., $= 1$) and then allow ρ to change. This is reasonable because in the former case, the leverage function $F_t^{\beta-1}$ can capture the dependence of volatility on the asset price (see Ren et al. 2007), while in the latter case, the negative correlation ρ can account for the leverage effect (see Hagan et al. 2002).

Our exact simulation method in these two cases possesses several appealing features.

(1) In contrast to the discretization schemes, the exact simulation method leads to unbiased estimates. This further results in a faster convergence rate for the total simulation errors.

(2) Compared to the existing asymptotic expansion methods for the SABR model, the errors of the exact simulation are easier to control because simply increasing the sample size N will make the errors converge to zero at a rate of $O(1/\sqrt{N})$.

(3) The exact simulation scheme is more flexible in several senses. First, it can be applied to pricing problems for not only European options but also other complicated path-dependent options such as the forward starting options. Second, it can be used for the SABR models with different behaviors at the boundary (the absorbing boundary and the reflecting boundary).

We plan to take three steps to achieve the exact simulation:

Step 1. Generate the terminal volatility.

Step 2. Simulate the integrated variance, given the terminal volatility.

Step 3. Sample the terminal price, given the terminal volatility and the integrated variance.

It turns out that our method performs accurately and efficiently. This is made possible because we overcome two main difficulties in the implementation about how to simulate two random variables whose distributions can be expressed in terms of the Hartman-Watson and the noncentral chi-squared distribution functions, respectively. We propose two novel simulation schemes to achieve numerical accuracy, efficiency, and stability. One stems from numerical Laplace inversion and Asian option literature, and the other is based on recent developments in evaluating the noncentral chi-squared distribution functions in a robust way. Please see Sections 3.2 and 3.3 for details.

When $\rho \neq 0$, our exact simulation method becomes a *semi-exact* one in that Steps 1 and 2 are exact, while only Step 3 is an approximation. Intuitively, the approximation is expected to be good when the time horizon T is small, and may become worse when T is big. This might be because the correlation between the price and volatility processes has a larger effect on the joint distribution at longer time intervals. Numerical results suggest that this semi-exact simulation method is still quite accurate when T is not big, e.g., no greater than one year.

However, when T becomes large, our semi-exact simulation method causes increasing biases. To reduce the biases, we propose a *piecewise semi-exact* (PSE) simulation method, where we first divide the time horizon $[0, T]$ into several short pieces with the same piece length Δ, and then apply the semi-exact simulation method piece by piece. It is worth pointing out that Δ needs not to be very small (i.e., the number of pieces needs not to be big). In fact, $\Delta = 0.5$ or 1 is usually short enough to make the PSE simulation method quite accurate. See Section 4.3.

This paper contributes to the literature in three ways. (i) It provides an alternative approach to tackling the computational issues related to the SABR model, one of the most popular stochastic volatility models in finance. In the two special but practically interesting cases, i.e., when the elasticity $\beta = 1$, or $\beta < 1$ and $\rho = 0$, our simulation is exact and has several advantages over the Euler scheme and existing asymptotic expansion approaches. When $\rho \neq 0$, the exact simulation method becomes a semi-exact one that is still quite accurate if the time horizon is not long. When the time horizon is long, a PSE simulation method is developed to reduce the biases significantly.

(ii) For European option pricing under the SABR model, we propose a novel, *conditional simulation* method, which reduces the variances of the plain simulation estimators substantially, e.g., by more than 99%. See Section 5 for more details. The idea of conditional simulation has been applied for European option pricing under the Heston model (see Broadie and Kaya 2006). A key step is that they express the conditional European option price, given the terminal volatility and the integrated variance, explicitly as a Black-Scholes formula, thanks to the conditional normal distribution of the asset price. However, in general cases, the conditional distribution of the forward price under the SABR model becomes much more complicated. We overcome this difficulty with the aid of the results by Carr and Linetsky (2006). See Theorem 5.1.

(iii) This paper also makes a contribution to the literature of the SDE simulation. Scholars recently began to pay attention to exact simulations of SDEs, stimulated by research developments in financial engineering. There have been two main threads in this direction. The first is to use the special structures of the processes in question. Broadie and Kaya (2004, 2006) exemplify the standard reference in this thread to discuss the exact simulation of affine stochastic volatility jump diffusion models. Giesecke et al. (2011) explore exact simulation for point processes with intensities driven by affine SDEs. Our algorithm also belongs to this category. But the nonlinearity of CEV diffusion in the price process excludes the SABR model from the affine processes, which makes this paper fundamentally different from these early studies. Chen et al. (2012) build up a simulation scheme of the SABR model based on the technique of moment matching. Inspired by Islah (2009) as well, they suggest to use lognormals to approximate the Hartman-Watson distribution with the moments of two distributions matched through small perturbation expansion. In comparison, their simulation is not exact even when the price and volatility processes are uncorrelated. Moreover, numerical results suggest that our

method is more efficient than theirs; see Sections 4.2.2 and 4.4.

The second thread of the SDE exact simulation is launched by Beskos and Roberts (2005). Their methodology combines the celebrated Girsanov Change of measure theorem with the acceptance-rejection sampling. Chen and Huang (2013) and Giesecke and Smelov (2013) extend this method to diffusions with unbounded drifts and jump diffusions, respectively. It is well known in the SDE literature (see, e.g., Chapter 5.5 of Karatzas and Shreve 1992) that *time change* and *removal of drift* are two major ways to solve an SDE. The literature started by Beskos and Roberts (2005) essentially attempts to remove the drift using the Girsanov theorem, while our paper aims at the time change technique. In this sense, our paper complements that literature too.

The remainder of this paper is organized as follows. In Section 2, we introduce the SABR model and the conditional distribution of the forward price given the volatility and integrated variance. Section 3 discusses implementation details of our exact simulation method. Numerical results are given in Section 4, where we also develop a PSE simulation scheme for long time horizons when the price and volatility processes are correlated. In Section 5, a conditional simulation method is developed for European option pricing. In Section 6, we extend our simulation method to the SABR model with a reflecting boundary. All technical proofs are deferred to the e-companion.

2. The SABR Model

The subsequent analysis uses a number of special functions and random variables. We collect them together here to make referencing easier. A *noncentral chi-squared* random variable $\chi'^2(\mu;\lambda)$ with μ degrees of freedom and noncentrality parameter λ has probability density function (pdf)

$$q_{\chi'^2}(x;\mu,\lambda)=\frac{1}{2}\exp\left(-\frac{x+\lambda}{2}\right)\left(\frac{x}{\lambda}\right)^{(\mu-2)/4} I_{\mu/2-1}(\sqrt{\lambda x})$$

for $x>0$, where $I_a(\cdot)$ is the *modified Bessel function of the first kind with index a* given by

$$I_a(x)=\sum_{k=0}^{+\infty}\frac{(x/2)^{a+2k}}{k!\Gamma(a+k+1)},\quad a\in\mathbb{R},$$

and $\Gamma(\cdot)$ is the gamma function. Denote its cumulative distribution function (cdf) by $Q_{\chi'^2}(x;\mu,\lambda)$. The *central chi-squared* random variable $\chi^2(\mu)$ is a special case of the noncentral one with $\lambda=0$. Its pdf $q_{\chi^2}(x;\mu)$ and cdf $Q_{\chi^2}(x;\mu)$ are given by

$$q_{\chi^2}(x;\mu)=\frac{e^{-x/2}x^{\mu/2-1}}{2^{\mu/2}\Gamma(\mu/2)}\quad\text{and}$$

$$Q_{\chi^2}(x;\mu)=\mathbb{P}[\chi^2(\mu)\leqslant x]=\int_0^x q_{\chi^2}(y;\mu)\,dy.$$

In addition, for any invertible function $h\colon\mathbb{R}\to\mathbb{R}$, we use h^{-1} to denote its inverse function.

2.1. The SABR Model

Consider a probability space $(\Omega,\mathcal{F},\mathbb{P})$, on which two independent standard Brownian motions $\{(W_t^{(1)},W_t^{(2)}):t\geqslant 0\}$ are defined. Assume that $\mathcal{G}^1$ and $\mathcal{G}^2$ are the natural filtrations generated by them, respectively, and let $\mathcal{F}=\mathcal{G}^1\otimes\mathcal{G}^2$. The SABR model describes the dynamics of an asset's forward price and its volatility, whose values at time t, $0\leqslant t\leqslant T$, are denoted by F_t and α_t, respectively. The SABR model is then given by the solution to the following SDEs:

$$dF_t=\alpha_t F_t^{\beta}\cdot[\sqrt{1-\rho^2}\,dW_t^{(1)}+\rho\,dW_t^{(2)}],\tag{1}$$

$$d\alpha_t=\nu\alpha_t\,dW_t^{(2)},\tag{2}$$

where the constants $\beta\in[0,1]$, $\rho\in(-1,1)$, and $\nu>0$. Apparently, the volatility α_t follows a geometric Brownian motion and F_t is governed by a CEV-type diffusion.

A subtle issue arises around the behavior of the forward price process $\{F_t\}$ at the boundary 0. We can show that F_t is always nonnegative for all $t\geqslant 0$. However, under some parameter ranges, it reaches 0 in finite time with a positive probability and the characterization of Equation (1) alone is not sufficient to determine the process uniquely. Additional specifications about the behavior of $\{F_t\}$ at 0 are hence needed. In this paper, we consider two ways to specify the boundary conditions: *absorbing* boundary and *reflecting* boundary. Roughly speaking, the former one forces the process to stay at 0 once it reaches the boundary; the latter is to "bump" $\{F_t\}$ back to the positive part immediately after it hits 0. More rigorous mathematical treatments on the boundary classification for a general SDE can be found in Borodin and Salminen (2002). In Sections 2–5, we focus on the absorbing boundary specification and leave the discussion of reflecting boundary to Section 6. Note that the absorbing boundary specification (for the CEV-type diffusions) is more valid for modeling the asset price than the reflecting boundary, because the latter causes an arbitrage opportunity. Indeed, one can buy the asset at zero cost at the moment it hits zero, and can earn a strictly positive return later due to the reflecting boundary. See, e.g., Andersen and Andreasen (2000), Davydov and Linetsky (2001), Henry-Labordère (2009), Doust (2012).

2.2. The Conditional Distribution of F_T Given α_0, α_T, $\int_0^T\alpha_s^2\,ds$, and F_0

In this paper, we intend to explore the possibility of exact simulation from the joint distribution of F_T and α_T. It turns out that our method relies on the conditional distribution of F_T given α_0, α_T, $\int_0^T\alpha_s^2\,ds$, and F_0.

The special case of $\beta = 1$ is relatively easy to treat. Given α_0, α_T, $\int_0^T \alpha_s^2\,ds$ and F_0, solving the SDE (1) and integrating the SDE (2) from 0 to T yields, respectively,

$$F_T = F_0 \exp\left\{-\frac{1}{2}\int_0^T \alpha_s^2\,ds + \rho\int_0^T \alpha_s\,dW_s^{(2)} + \sqrt{1-\rho^2}\int_0^T \alpha_s\,dW_s^{(1)}\right\}, \tag{3}$$

$$\int_0^T \alpha_s\,dW_s^{(2)} = \frac{\alpha_T - \alpha_0}{\nu}. \tag{4}$$

Substituting (4) into (3), we obtain

$$F_T = F_0 \exp\left\{-\frac{1}{2}\int_0^T \alpha_s^2\,ds + \frac{\rho}{\nu}(\alpha_T - \alpha_0) + \sqrt{1-\rho^2}\int_0^T \alpha_s\,dW_s^{(1)}\right\}.$$

Since $\{W_s^{(1)}: 0 \leqslant s \leqslant T\}$ and $\{\alpha_s: 0 \leqslant s \leqslant T\}$ are independent, it follows that provided α_0, α_T, $\int_0^T \alpha_s^2\,ds$, and F_0, the conditional distribution of $\ln F_T$ is normal with mean $\ln F_0 - \frac{1}{2}\int_0^T \alpha_s^2\,ds + (\rho/\nu)(\alpha_T - \alpha_0)$ and variance $(1-\rho^2)\int_0^T \alpha_s^2\,ds$.

Unfortunately, in the general case of $\beta < 1$, it seems very difficult to derive the exact joint distribution of F_T and α_T. To the authors' best knowledge, there exists no "complete" answer to this question in the literature. A nice "partial" solution is found by Islah (2009) and presented in Propositions 2.1 and 6.1. See Section EC.2 in the e-companion for an outline of the related proofs.

Proposition 2.1 (Islah 2009). *Fix $T > 0$ and suppose that α_0, α_T, $\int_0^T \alpha_s^2\,ds$, and F_0 are given.*

(i) *When $\beta = 1$, the conditional distribution of F_t is log-normal and satisfies*

$$\ln F_T \sim N\left(\ln F_0 - \frac{1}{2}\int_0^T \alpha_s^2\,ds + \frac{\rho}{\nu}(\alpha_T - \alpha_0),\ (1-\rho^2)\int_0^T \alpha_s^2\,ds\right);$$

(ii) *When $\beta \in [0,1)$, $\rho = 0$, and $\{F_t\}$ have an absorbing boundary at 0, the conditional distribution of F_t can be expressed in terms of noncentral chi-squared distribution functions, i.e.,*

$$\mathbb{P}\left(F_T = 0 \,\middle|\, F_0, \alpha_0, \alpha_T, \int_0^T \alpha_s^2\,ds\right) = 1 - Q_{\chi^2}\left(A_0; \frac{1}{1-\beta}\right); \tag{5}$$

$$\mathbb{P}\left(F_T \leqslant u \,\middle|\, F_0, \alpha_0, \alpha_T, \int_0^T \alpha_s^2\,ds\right) = 1 - Q_{\chi'^2}\left(A_0; \frac{1}{1-\beta}, C_0(u)\right) \quad \text{for any } u > 0, \tag{6}$$

where

$$A_0 = \frac{1}{\int_0^T \alpha_s^2\,ds}\left(\frac{F_0^{1-\beta}}{1-\beta}\right)^2 \quad \text{and} \quad C_0(u) = \frac{1}{\int_0^T \alpha_s^2\,ds}\cdot\frac{u^{2(1-\beta)}}{(1-\beta)^2};$$

(iii) *When $\beta \in [0,1)$, $\rho \neq 0$, and $\{F_t\}$ have an absorbing boundary at 0, the conditional distribution of F_T can be approximated by the following*:

$$\mathbb{P}\left(F_T = 0 \,\middle|\, F_0, \alpha_0, \alpha_T, \int_0^T \alpha_s^2\,ds\right) \approx 1 - Q_{\chi^2}\left(A; 1 + \frac{\beta}{(1-\beta)(1-\rho^2)}\right); \tag{7}$$

$$\mathbb{P}\left(F_T \leqslant u \,\middle|\, F_0, \alpha_0, \alpha_T, \int_0^T \alpha_s^2\,ds\right) \approx 1 - Q_{\chi'^2}\left(A; 1 + \frac{\beta}{(1-\beta)(1-\rho^2)}, C(u)\right) \quad \text{for any } u > 0, \tag{8}$$

where

$$A = \frac{1}{(1-\rho^2)\int_0^T \alpha_s^2\,ds}\left(\frac{F_0^{1-\beta}}{1-\beta} + \frac{\rho}{\nu}(\alpha_T - \alpha_0)\right)^2 \quad \text{and} \quad C(u) = \frac{1}{(1-\rho^2)\int_0^T \alpha_s^2\,ds}\cdot\frac{u^{2(1-\beta)}}{(1-\beta)^2}.$$

Remark 2.1. When $\beta = 1$, or $\beta \in [0,1)$ and $\rho = 0$, Proposition 2.1 provides an exact conditional distribution of F_T given α_0, α_T, $\int_0^T \alpha_u^2\,du$, and F_0. In Section 3, we shall propose an *exact* simulation method to sequentially simulate α_T, $\int_0^T \alpha_u^2\,du$, and F_T. Note that these two special cases are also interesting and important for the SABR model in practice. For more details, please see the seventh paragraph in Section 1.

Remark 2.2. When $\beta \in [0,1)$ and $\rho \neq 0$, (7)–(8) serve as an approximation to the conditional cdf of F_T. Therefore the exact simulation method for $\rho = 0$ in Section 3 becomes a *semi-exact* one because the simulations of α_T and $\int_0^T \alpha_u^2\,du$ are still exact, whereas the simulation of F_T given α_T and $\int_0^T \alpha_u^2\,du$ is based on the approximate conditional cdf (7)–(8). However, numerical results suggest that this semi-exact simulation method is still quite accurate when T is not big, e.g., $T \leqslant 1$. When T is big, we will propose a PSE simulation scheme to reduce the biases substantially. Specifically, we first divide the time horizon into a small number of pieces, e.g., five pieces when $T = 5$, and then employ the semi-exact simulation piece by piece. See Section 4.3 for a detailed discussion.

Remark 2.3. When $\beta \in [0,1)$, $\rho = 0$, and 0 is specified as an absorbing boundary, the conditional distribution

F_T has a positive probability mass at the point 0. Moreover, it is worth pointing out that on the right-hand side (RHS) of Equation (6), the argument u appears not in A_0 but in $C_0(u)$. This means that F_T is not a noncentral chi-squared random variable. We only know that its cdf can be *expressed in terms of* the noncentral chi-squared cdf. Therefore, existing algorithms of simulating noncentral chi-squared random variables no longer apply. We resolve this issue by employing the inverse transform method. Similar difficulty is also encountered for the semi-exact simulation in the case of $\rho \neq 0$. See Section 3.3 for details.

3. Simulation of the SABR Model

Based on Proposition 2.1 and Remarks 2.1–2.2, we propose a simulation scheme for the SABR model to generate sample pairs of F_T and α_T, given F_0 and α_0. It is divided into three steps.

A Simulation Method for the SABR Model

Step 1. Generate a sample from the distribution of α_T, given α_0.

Step 2. Sample from $\int_0^T \alpha_s^2\, ds$, given α_0 and α_T.

Step 3. Generate a sample from the distribution of F_T, given α_0, α_T, $\int_0^T \alpha_s^2\, ds$, and F_0.

3.1. Sampling from α_T Given α_0

Since the volatility is driven by a geometric Brownian motion, it follows that

$$\alpha_T = \alpha_0 \exp\left[-\frac{1}{2}\nu^2 T + \nu W_T^{(2)}\right].$$

Therefore the conditional distribution of $\ln \alpha_T$, given α_0, is normal with mean $\ln \alpha_0 - \frac{1}{2}\nu^2 T$ and variance $\nu^2 T$. Then, it is straightforward to sample from α_T.

3.2. Sampling from $\int_0^T \alpha_s^2\, ds$ Given α_0 and α_T

To sample from $\int_0^T \alpha_s^2 ds$ given α_0 and α_T, the conditional distribution of $\int_0^T \alpha_s^2\, ds$ apparently plays a crucial role. Yor (1980) derives an explicit expression for the conditional pdf of a similar r.v.,

$$A_t^{(\mu)} := \int_0^t \exp(2B_s^{(\mu)})\, ds, \tag{9}$$

where $B_t^{(\mu)} := B_t + \mu t$ with $\{B_t\colon t \geqslant 0\}$ being a standard Brownian motion and μ a constant; see also, e.g., Equation (1) on p. 1051 in Barrieu et al. (2004). Applying Yor's result, we obtain Proposition 3.1.

Proposition 3.1. *The density function of $\int_0^T \alpha_s^2\, ds$, conditioning on α_0 and α_T, is given by*

$$P\left(\int_0^T \alpha_s^2\, ds \in dw \,\middle|\, \alpha_0, \alpha_T\right)$$

$$= \frac{\sqrt{2\pi T}\nu}{w} \exp\left\{\frac{1}{2\nu^2}\left(\frac{1}{T}\left[\ln\left(\frac{\alpha_T}{\alpha_0}\right)\right]^2 - \frac{\alpha_0^2 + \alpha_T^2}{w}\right)\right\} \cdot I_0(\hat{r}) f_{\hat{r}}(\nu^2 T)\, dw, \tag{10}$$

where $\hat{r} \equiv \hat{r}(w) := \alpha_0\alpha_T/(\nu^2 w)$, $I_a(\cdot)$ for $a \in \mathbb{R}$ denotes the modified Bessel function of the first kind with index a, and $f_r(\cdot)$ for $r > 0$ is the Hartman-Watson density function

$$f_r(t) = \frac{1}{I_0(r)} \frac{r}{\sqrt{2\pi^3 t}} \exp\left(\frac{\pi^2}{2t}\right) \int_0^{+\infty} \exp\left(-\frac{y^2}{2t}\right) \cdot \exp(-r\cosh(y)) \sinh(y) \sin\left(\frac{\pi y}{t}\right) dy, \quad t > 0. \tag{11}$$

Provided the explicit expression (10) of the conditional pdf, it is tempting to integrate it numerically to obtain the conditional cdf and generate a sample of $\int_0^T \alpha_s^2\, ds$, by the acceptance-rejection method or by the inverse transform method. In either case, we need to repeatedly evaluate the conditional pdf (10), and thereby the Hartman-Watson density function $f_r(\cdot)$ in (11). However, as observed by Yor (1992), Barrieu et al. (2004), Boyle and Potapchik (2006), among others, $f_r(t)$ is extremely difficult to compute for small t. Indeed, as t tends to 0, the component $(r/\sqrt{2\pi^3 t})\exp(\pi^2/(2t))$ outside the integral in (11) grows exponentially to infinity, while the sign of the component $\sin(\pi y/t)$ inside the integral changes more and more frequently. If high-precision computations are not employed, this would lead to high inaccuracy for the numerical outcomes of $f_r(t)$; see Figures 1–3 in Barrieu et al. (2004) for this phenomenon. To achieve an acceptable result, Boyle and Potapchik (2006) point out that one has to compute the integral in (11) *highly precisely*. They use a 90-digit precision during the computation to generate adequately accurate approximations to $f_{0.5}(t)$ when $t = 0.1$. Nonetheless, this is not appealing from the perspective of algorithm efficiency.

To overcome this difficulty, we intend to evaluate the conditional cdf by inverting its Laplace transform rather than integrating (10) numerically, and then exploit the inverse transform method to generate a sample of $\int_0^T \alpha_s^2\, ds$. In the interest of algorithm efficiency, we surely prefer a simpler form of the Laplace transform, which is not available for $\int_0^T \alpha_s^2\, ds$ itself, unfortunately; see Equation (5.5) in Matsumoto and Yor (2005) for a complicated expression in terms of a one-dimensional integral with the integrand involving the Bessel function of the first kind.

However, we find that this idea is still applicable if there exists a bijection $h(\cdot)$, defined in $\mathbb{R}^+$, such that a simple closed-form Laplace transform of the conditional cdf of $h(\int_0^T \alpha_s^2\, ds)$ is available. Denote by $L_h(\cdot)$

the cdf of $h(\int_t^T \alpha_s^2\, ds)$ given α_0 and α_T, and by $\hat{L}_h(\cdot)$ its Laplace transform, i.e.,

$$\hat{L}_h(\theta) := \int_0^{+\infty} e^{-\theta u} L_h(u)\, du, \quad \text{for } \theta > 0.$$

If $\hat{L}_h(\cdot)$ is known explicitly, we can evaluate $L_h(\cdot)$ numerically through some Laplace inversion methods. Then, given a sample $U \sim \text{Unif}[0,1]$, applying root-finding algorithms can solve the equation $L_h(V) = U$ numerically to obtain the root V. Since V follows the same distribution as $h(\int_0^T \alpha_s^2\, ds)$ conditioning on α_0 and α_T, then $h^{-1}(V)$ is a sample of $\int_0^T \alpha_s^2\, ds$ given α_0 and α_T.

Consider a special bijection $h(x) = 1/x$ for $x > 0$. Matsumoto and Yor (2005) show that

$$E\left[\exp\left(-\frac{\lambda}{A_t^{(\mu)}}\right)\middle| B_t^{(\mu)} = x\right] = \exp\left(-\frac{\phi_x(\lambda)^2 - x^2}{2t}\right) \quad \text{for any } t > 0 \text{ and } \lambda > 0, \tag{12}$$

where $A_t^{(\mu)}$ is defined in (9) and

$$\phi_x(\lambda) := \operatorname{arcosh}(\lambda e^{-x} + \cosh(x)). \tag{13}$$

Here, $\cosh(\cdot)$ and $\operatorname{arcosh}(\cdot)$ denote the *hyperbolic cosine* and *inverse hyperbolic cosine* functions, respectively, i.e.,

$$\cosh(y) = \frac{e^y + e^{-y}}{2} \quad \text{and} \quad \operatorname{arcosh}(z) = \ln(z + \sqrt{z^2 - 1}). \tag{14}$$

Based on (12), we establish the following proposition.

Proposition 3.2. *When $h(x) = 1/x$, the Laplace transform of $L_h(\cdot)$ is given by*

$$\hat{L}_h(\theta) = \frac{1}{\theta} \exp\left\{-\frac{[\phi_{\ln(\alpha_T/\alpha_0)}(\theta\nu^2/\alpha_0^2)]^2 - [\ln(\alpha_T/\alpha_0)]^2}{2\nu^2 T}\right\}, \tag{15}$$

where $\phi_x(\lambda)$ is defined in (13).

Remark 3.1. Note that $\hat{L}_h(\theta)$ with $h(x) = 1/x$ has a surprisingly simple form (15) that involves elementary functions only. This appealing feature enables us to evaluate $L_h(u)$ very efficiently and accurately through Laplace inversion. To the authors' best knowledge, most of the exact simulation methods resorting to the Laplace transform to generate SDE samples choose $h(x) = x$ (see, e.g., Broadie and Kaya 2006, Giesecke et al. 2011). In contrast, our paper selects $h(x) = 1/x$.

There has been extensive research on numerical inversion of Laplace transforms. Here, we make use of the Euler inversion algorithm proposed by Abate and Whitt (1992), which has been applied widely in financial engineering, as well as general operations research due to its high efficiency and accuracy. In addition, the Euler inversion formula is very simple and thus easy to implement. Specifically, given the Laplace transform $\hat{L}_h(\theta)$, the original function $L_h(u)$ reads

$$L_h(u) = \frac{e^{M/2}}{2u}\operatorname{Re}\left(\hat{L}_h\left(\frac{M}{2u}\right)\right) + \frac{e^{M/2}}{u}\sum_{k=1}^{+\infty}(-1)^k \operatorname{Re}\left(\hat{L}_h\left(\frac{M - 2k\pi i}{2u}\right)\right) - e_d, \tag{16}$$

where $\operatorname{Re}(x)$ is the real part of x, M a positive constant, $i = \sqrt{-1}$, and e_d the discretization error

$$e_d \equiv e_d(L_h, u, M) := \sum_{k=1}^{+\infty} e^{-kM} L_h((2k+1)u).$$

One advantage of the Euler inversion algorithm is that we can easily control the discretization error by choosing a sufficiently large M. It follows immediately from $0 \leqslant L_h \leqslant 1$ that

$$0 \leqslant e_d \leqslant \frac{e^{-M}}{1 - e^{-M}} \approx e^{-M} \quad \text{for sufficiently large } M. \tag{17}$$

In our implementation, M is chosen to be 20 as suggested by Abate and Whitt (1992).

Besides, when computing the (nearly) alternating series of the form $\sum_{k=0}^{+\infty}(-1)^k a_k$ in (16), we accelerate the convergence via Euler transformation, i.e., we approximate $\sum_{k=0}^{+\infty}(-1)^k a_k$ by

$$E(m, n) := \sum_{k=0}^{m} \frac{m!}{k!(m-k)!} 2^{-m} S_{n+k},$$

where $S_j := \sum_{k=0}^{j}(-1)^k a_k$. We set two parameters $m = 20$ and $n = 35$ in our implementation. See Abate and Whitt (1992) for more details about the Euler inversion method and Euler transformation.

Remark 3.2. When T is small, the exponent on the RHS of (15) could be large, either positive or negative. We study this effect on the accuracy of the Euler inversion algorithm as well as the selection of involved algorithm parameters m and n in Section EC.5 in the e-companion. We suggest a test to make a proper choice of m and n to achieve a desired accuracy. It turns out that when T is no less than one month, which is typical for common financial applications, choosing $m = 20$ and $n = 35$ has achieved a high accuracy. See Section EC.5 for details.

In summary, we go through the following two substeps to complete Step 2, namely, to generate a sample

of $\int_0^T \alpha_s^2\,ds$, given α_0 and α_T, from its exact conditional distribution.

Two Substeps in Step 2

Substep 2.1. Generate a sample of $V = h(\int_0^T \alpha_s^2\,ds) = 1/\int_0^T \alpha_s^2\,ds$, given α_0 and α_T, via the inverse transform method.

Substep 2.2. Set $\int_0^T \alpha_s^2\,ds = 1/V$, i.e., $1/V$ is a sample of $\int_0^T \alpha_s^2\,ds$ given α_0 and α_T.

In Substep 2.1, we need to solve the equation $L_h(V) = U$ to find V. Due to the monotonicity of L_h, simple root-finding algorithms such as Newton's method and the bisection method are applicable. At each iteration of the search for the root, the value of L_h can be calculated quickly and accurately by inverting its Laplace transform $\hat{L}_h(\cdot)$ numerically via the Euler inversion algorithm. The numerical tolerance of the search is set to be 10^{-6} to achieve the desired accuracy in our numerical experiments.

3.3. Sampling from F_T Given α_0, α_T, $\int_0^T \alpha_s^2\,ds$, and F_0

When $\beta = 1$, the conditional distribution of F_T given α_T and $\int_0^T \alpha_s^2\,ds$ is log-normal by Proposition 2.1. Therefore, it is simple to sample from F_T.

When $0 \leqslant \beta < 1$, we shall focus only on the case of $\rho = 0$, where an exact simulation method will be proposed to sample from F_T given α_0, α_T, $\int_0^T \alpha_s^2\,ds$, and F_0. When $\rho \neq 0$, a similar method will serve as an approximation. As discussed in Remark 2.3, when $\rho = 0$, the conditional distribution of F_T given in Proposition 2.1 is *not* a noncentral chi-squared distribution. Therefore the techniques of simulating the noncentral chi-squared random variables, such as the ones used in Broadie and Kaya (2006), no longer apply in our case. We shall employ the inverse transform method as follows to achieve the purpose. Define P_0 to be the conditional probability of F_T at 0 given by (5). Then, we first generate a sample $U \sim \text{Unif}[0,1]$. If $U \leqslant P_0$, then set $F_T = 0$. Otherwise, we use the root-finding algorithm to find the root $\hat{U}$ to the following equation, and then set $F_T = \hat{U}$:

$$\mathbb{P}\left(F_T \leqslant \hat{U} \,\middle|\, F_0, \alpha_0, \alpha_T, \int_0^T \alpha_s^2\,ds\right) = 1 - Q_{\chi'^2}\left(A_0; \frac{1}{1-\beta}, C_0(\hat{U})\right) = U. \tag{18}$$

When searching the root of (18), we are required to evaluate the noncentral chi-squared cdf $Q_{\chi'^2}(A_0; B, C_0(\hat{U}))$ repeatedly for fixed A_0 and $B \equiv 1/(1-\beta)$, but various values of $C_0(\hat{U})$. Moreover, in a complete run of our exact sampling of (F_T, α_T), although the value of B is always fixed by the model setting, the values of A_0 and $C_0(\hat{U})$ vary with $\int_0^T \alpha_s^2\,ds$, and $\int_0^T \alpha_s^2\,ds$ and $\hat{U}$,

Figure 1. (Color online) The 1,259 Pairs of $(A_0, C_0(\hat{U}))$ of $Q_{\chi'^2}(A_0; B, C_0(\hat{U}))$ That We Encounter and Hence Need to Evaluate When Exactly Generating 100 Samples from (F_T, α_T)

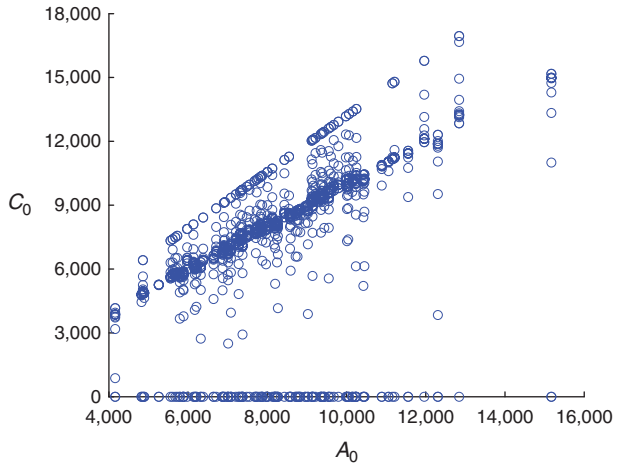

Notes. Associated parameters are $\nu = 0.3$, $\beta = 0.8$, $\rho = 0$, $T = 0.5$, $\alpha_0 = 0.2$, and $F_0 = 100$. We can see that A_0 varies over [4,161, 15,172] and $C_0(\hat{U})$ over [0, 16,948].

respectively. It can be seen from their expressions that they are very sensitive to changes of these values under many circumstances, for example, when α_0 is small or β is close to 1. As a result, A_0 and $C_0(\hat{U})$ may vary over very broad ranges. Figure 1 shows all the pairs of $(A_0, C_0(\hat{U}))$ of $Q_{\chi'^2}(A_0; B, C_0(\hat{U}))$ that we encounter, and thereby need to evaluate when generating 100 samples from (F_T, α_T) with $\alpha_0 = 0.2$ and $\beta = 0.8$. We can see that A_0 varies over [4,161, 15,172] and $C_0(\hat{U})$ over [0, 16,948].

It is clear from the discussion above that an algorithm capable of evaluating the noncentral chi-squared cdf $Q_{\chi'^2}(x; \delta, \lambda)$ efficiently and accurately for very broad ranges of parameters is crucial to our simulation. Although a variety of algorithms have been designed for computing $Q_{\chi'^2}(x; \delta, \lambda)$ (see, e.g., Sankaran 1963, Schroder 1989, Ding 1992, Penev and Raykov 2000), most of them perform well only if the parameters vary over certain narrow ranges, and always become much less efficient otherwise. Recently, Larguinho et al. (2011) proposed a two-part strategy that combines the recursive method by Ding (1992) with the analytical approximation by Penev and Raykov (2000). Numerical experiments demonstrate that this strategy is consistently efficient for all the parameters x and λ, while guaranteeing high accuracy with the absolute error no greater than 10^{-6}.

To be specific, when the parameters $x < 500$ and $\lambda < 500$, we adopt the following recursive algorithm proposed by Ding (1992):

$$Q_{\chi'^2}(x; \delta, \lambda) = \sum_{k=0}^{+\infty} y_k t_k, \tag{19}$$

where

$$t_0 = \frac{1}{\Gamma(\delta/2+1)}\left(\frac{x}{2}\right)^{\delta/2} e^{-x/2}, \quad y_0 = u_0 = e^{-\lambda/2},$$
$$t_k = t_{k-1}\frac{x}{\delta+2k}, \quad y_k = y_{k-1} + u_k,$$
$$u_k = \frac{u_{k-1}\lambda}{2k} \quad \text{for } k \geqslant 1.$$

Ding (1992) argues that when $\delta + 2n > x$, the truncation error of the series (19) can be bounded by $\sum_{k=n}^{+\infty} y_k t_k \leqslant t_{n-1}x/(\delta+2n-x)$. In our implementation, we select n such that the truncation error $\leqslant 10^{-7}$.

In contrast, when the parameters $x \geqslant 500$ or $\lambda \geqslant 500$, we apply the analytic approximation of Penev and Raykov (2000), $Q_{\chi'^2}(x;\delta,\lambda) \approx \Psi(z)$, where $\Psi(\cdot)$ denotes the standard normal cdf and

$$z = \operatorname{sign}(s-1)\cdot\left[\delta\cdot(s-1)^2\left(\frac{1}{2s}+\mu^2-\frac{\kappa(1-s)}{s}\right) - \ln\left(\frac{1}{s}-\frac{2\kappa(1-s)}{s(1+2\mu^2 s)}\right)+\frac{2\Theta(s)}{\delta}\right]^{1/2},$$

with

$$\mu^2 = \lambda/\delta, \quad s = \frac{\sqrt{1+4x\mu^2/\delta}-1}{2\mu^2},$$
$$\eta = \frac{1+2\mu^2 s-2\kappa(1-s)-s-2\mu^2 s^2}{1+2\mu^2 s-2\kappa(1-s)},$$
$$\kappa(s) = \frac{(1-s)\ln(1-s)+s-s^2/2}{s^2},$$
$$\Theta(s) = -\frac{3}{2}\frac{1+4\mu^2 s}{(1+2\mu^2 s)^2}+\frac{5}{3}\frac{(1+3\mu^2 s)^2}{(1+2\mu^2 s)^3} + \frac{2(1+3\mu^2 s)}{(s-1)(1+2\mu^2 s)^2}+\frac{3\eta}{(s-1)^2(1+2\mu^2 s)} - \frac{(1+2\kappa(\eta))\eta^2}{2(s-1)^2(1+2\mu^2 s)}.$$

Although no analytical result about the bound of the approximation error exists in Penev and Raykov (2000) and subsequent literature, Dyrting (2004) reports that this method can guarantee an accuracy of at least 10^{-6} over the proposed parameter ranges by some numerical experiments.

To summarize, when $0 \leqslant \beta < 1$, we take the following two substeps to complete Step 3.

Two Substeps in Step 3

Substep 3.1. Generate a uniform random variable $U \sim \text{Unif}(0,1)$. If $U \leqslant P_0$, then set $F_T = 0$ and the algorithm terminates. Otherwise, go to Substep 3.2.

Substep 3.2. Sample from F_T, given F_0, α_0, α_T, and $\int_0^T \alpha_s^2\,ds$, via the inverse transform method. Specifically, use the root-finding algorithm to find the root $\hat{U}$ of the equation $\mathbb{P}(F_T \leqslant \hat{U} \mid F_0, \alpha_0, \alpha_T, \int_0^T \alpha_s^2\,ds) = U$, and then set $F_T = \hat{U}$.

4. Numerical Results and a PSE Simulation Scheme

In Subsections 4.1–4.2, we provide numerical comparisons between the proposed exact simulation method for $\rho = 0$ and other existing numerical methods for the SABR model. Our method appears to be very accurate. The semi-exact simulation method for $\rho \neq 0$ is investigated in Subsection 4.3. Numerical examples suggest that this semi-exact simulation method is still quite accurate when T is not big, e.g., $T \leqslant 1$. When T is large, we propose a PSE simulation method to improve the accuracy substantially. Subsection 4.4 compares numerically the PSE simulation method with the Euler scheme and the low-bias scheme of Chen et al. (2012). Subsection 4.5 gives an example of pricing exotic options through our simulation method. All the computations are completed using Matlab 7 on a desktop with an Intel Core2 Q9400 2.66 GHZ processor.

4.1. Test of the Accuracy of the Exact Simulation Method

This subsection is devoted to testing the accuracy of the exact simulation method for $\rho = 0$ using European call option prices. Here, we apply the finite difference method (FDM) to the associated PDEs for the call option prices and use the resulting values as benchmarks, where five algorithm parameters are involved, two for the computing domains and three others for the discretization stepsizes. To accelerate the convergence, we use Yanenko's scheme with Richardson extrapolation (see, e.g., Duffy 2006).

First, we consider three parameter settings in Table 1, corresponding to the interest rate market ($F_0 = 0.05$), the foreign exchange market ($F_0 = 1.10$), and the equity derivatives market ($F_0 = 100$). In these settings, the expansion formula by Hagan et al. (2002) is expected to be highly accurate because T, $\alpha_0 F_0^{\beta-1}\nu$ and $\log(K/F_0)$ are all small. Note that the corrected version of the pricing formula of Hagan et al. (2002) by Obłój (2008) turns out to generate the same results as the former.

Table 2 gives European call option prices obtained by the exact simulation method, the expansion formula of Hagan et al. (2002), and the FDM. Note that we simulate a huge number of sample paths (10,240,000) to generate the exact simulation prices. This results in very small standard errors. For instance, the standard errors in Case I.A are at the level of 10^{-6}. However, all the benchmarks still lie in the quite narrow confidence

Table 1. Parameter Settings for Testing the Accuracy of the Exact Simulation Method

Case number	F_0	Range of K	T	ρ	β	ν	α_0
I.A (Interest rate market)	0.05	[0.045, 0.055]	1	0	0.55	0.03	0.20
I.B (Foreign exchange market)	1.10	[1.00, 1.20]	1	0	0.70	0.10	0.20
I.C (Equity derivatives market)	100	[90, 110]	1	0	0.60	0.20	0.30

intervals (CIs) of our exact simulation estimates. This indicates that our exact simulation method is highly accurate and reliable.

Second, we attempt to test the accuracy of our exact simulation method for $\rho = 0$ when the expansion formula of Hagan et al. (2002) is not reliable. To this end, we consider three parameter settings in Table 3 in the foreign exchange market, where we let the maturity T, the volatility of volatility ν, and the strike K vary in wide ranges, respectively. The numerical results are demonstrated in Figure 2. In Cases II.A and II.B (the left and middle panels), when T is bigger than 4 (years) or ν is bigger than 0.5, the expansion formula of Hagan et al. (2002) incurs an increasing upward bias. However, the results of our exact simulation method always coincide with those of the FDM. This is not surprising because as mentioned earlier, Hagan's expansion formula may become less accurate when either T or $\alpha_0 F_0^{\beta-1}\nu$ is large,

Table 2. Testing the Accuracy of the Exact Simulation Method When the Expansion Formula by Hagan et al. (2002) Is Highly Accurate

K	Exact simulation	Standard errors	95% confidence intervals	Hagan et al.	FDM
Prices of European call options in the interest rate market (Case I.A)					
0.045	0.01727	9.19E–06	[0.01725, 0.01728]	0.01726	0.01725
0.046	0.01681	9.09E–06	[0.01679, 0.01682]	0.01680	0.01679
0.047	0.01636	8.99E–06	[0.01634, 0.01637]	0.01635	0.01634
0.048	0.01592	8.89E–06	[0.01590, 0.01593]	0.01591	0.01590
0.049	0.01549	8.79E–06	[0.01547, 0.01550]	0.01548	0.01547
0.050	0.01507	8.69E–06	[0.01505 , 0.01509]	0.01506	0.01505
0.051	0.01466	8.60E–06	[0.01464, 0.01468]	0.01465	0.01464
0.052	0.01426	8.50E–06	[0.01424, 0.01427]	0.01425	0.01424
0.053	0.01387	8.40E–06	[0.01385, 0.01388]	0.01386	0.01385
0.054	0.01349	8.30E–06	[0.01347, 0.01350]	0.01348	0.01347
0.055	0.01311	8.20E–06	[0.01310, 0.01313]	0.01311	0.01310
Prices of European call options in the foreign exchange market (Case I.B)					
1.00	0.14188	5.21E–05	[0.14178, 0.14198]	0.14196	0.14197
1.02	0.12903	5.03E–05	[0.12893, 0.12913]	0.12911	0.12911
1.04	0.11693	4.84E–05	[0.11684, 0.11703]	0.11701	0.11700
1.06	0.10559	4.64E–05	[0.10550, 0.10568]	0.10566	0.10565
1.08	0.09500	4.44E–05	[0.09492, 0.09509]	0.09507	0.09507
1.10	0.08518	4.23E–05	[0.08509, 0.08526]	0.08524	0.08523
1.12	0.07609	4.02E–05	[0.07601, 0.07617]	0.07616	0.07615
1.14	0.06774	3.81E–05	[0.06766, 0.06781]	0.06780	0.06779
1.16	0.06009	3.61E–05	[0.06002, 0.06016]	0.06014	0.06013
1.18	0.05312	3.40E–05	[0.05305, 0.05318]	0.05316	0.05316
1.20	0.04679	3.20E–05	[0.04673, 0.04685]	0.04683	0.04683
Prices of European call options in the equity derivatives market (Case I.C)					
90	10.02941	1.48E–03	[10.02652, 10.03230]	10.03079	10.03078
92	8.08810	1.44E–03	[8.08527, 8.09092]	8.08957	8.08958
94	6.22929	1.37E–03	[6.22660, 6.23198]	6.23083	6.23085
96	4.52389	1.26E–03	[4.52143, 4.52636]	4.52561	4.52561
98	3.05762	1.09E–03	[3.05548, 3.05977]	3.05943	3.05938
100	1.90169	8.94E–04	[1.89994, 1.90344]	1.90301	1.90294
102	1.07945	6.83E–04	[1.07811, 1.08079]	1.08033	1.08027
104	0.55708	4.89E–04	[0.55612, 0.55803]	0.55771	0.55769
106	0.26164	3.29E–04	[0.26100, 0.26229]	0.26209	0.26209
108	0.11232	2.11E–04	[0.11191, 0.11273]	0.11269	0.11269
110	0.04444	1.30E–04	[0.04419, 0.04469]	0.04469	0.04468

Notes. We can see that the two benchmarks of Hagan et al. and the FDM lie in the very narrow CIs of our exact simulation estimates. This suggests that our method is highly accurate.

Table 3. Parameter Settings for Testing the Accuracy of the Exact Simulation Method

Case number	F_0	K	T	ρ	β	ν	α_0
II.A (Maturity varies)	1.10	1.10	[1, 10]	0	0.80	0.40	0.30
II.B (Vol. of vol. varies)	1.10	1.10	1	0	0.80	[0.05, 0.95]	0.30
II.C (Strike varies)	1.10	[0.50, 2.30]	1	0	0.80	0.40	0.30

but the exact simulation method will consistently produce the unbiased estimator. In Case II.C (the right panel), all three methods generate almost identical numerical prices even when the European options are deep in the money ($K/F_0 = 50\%$) and deep out of money ($K/F_0 = 200\%$). The pricing errors are not visible in this case.

4.2. Comparison between the Exact Simulation and Other Simulation Schemes

4.2.1. Comparison with the Euler Scheme. Broadie and Kaya (2006) use the European call option to conduct a similar comparison between their exact simulation and the Euler scheme under the Heston model. Their numerical results indicate that the exact simulation is more accurate than the Euler scheme, because the former is unbiased, but the latter incurs a bias due to the Euler discretization, and moreover, the former can lead to a faster convergence rate. We shall illustrate that, not surprisingly, our exact simulation method of the SABR model for $\rho = 0$ possesses similar advantages over the Euler scheme.

Intuitively, when the "real" initial volatility $\alpha_0 F_0^{\beta-1}$ and the volatility of volatility ν get larger, the forward price $\{F_t\colon t \in [0,T]\}$ will become more volatile. This would then lead to more distortions of the distribution of F_T for the Euler scheme, and hence, a larger bias. To reduce the bias, one has to increase the number of discretization time steps at the cost of computation time.

Consider the parameter settings with relatively large $\alpha_0 F_0^{\beta-1}$ and ν and with different maturities (one year, threeyears and five years) in Table 4. We shall compare the root mean square (RMS) errors between the exact simulation method and the Euler scheme. In the Euler scheme, once a negative value is generated in an intermediate step, we fix its value to be zero until maturity.

For a simulation estimator $\widehat{\text{EuP}}$ of the true European option price EuP, its RMS error is defined as

$$\text{RMS error} := \sqrt{\text{Bias}^2 + \text{Variance}},$$

where

$$\text{Bias} := \mathbb{E}(\widehat{\text{EuP}}) - \text{EuP} \quad \text{and}$$
$$\text{Variance} := \mathbb{E}[(\widehat{\text{EuP}} - \mathbb{E}(\widehat{\text{EuP}}))^2].$$

For a fixed sample size N and a fixed number of time steps M, denote by $\widehat{\text{EuP}}_1(N,M)$ and $\widehat{\text{EuP}}_2(N)$ the Euler simulation estimator and the exact simulation estimator for the European option price under the SABR model, respectively. Since $\widehat{\text{EuP}}_2(N)$ is unbiased when $\rho = 0$, the corresponding RMS error is simply given by its standard error. By contrast, the RMS error of $\widehat{\text{EuP}}_1(N,M)$ consists of the standard error and a nonzero bias. Besides, we shall compute the true option price EuP via the FDM, and estimate the mean of the Euler simulation estimator $\mathbb{E}[\widehat{\text{EuP}}_1(N,M)]$ using $\widehat{\text{EuP}}_1(N^*,M)$ with N^* being a very large number such as 10,240,000.

Figure 2. (Color online) Testing the Accuracy of the Exact Simulation Method When the Expansion Formula by Hagan et al. (2002) Is Not Reliable

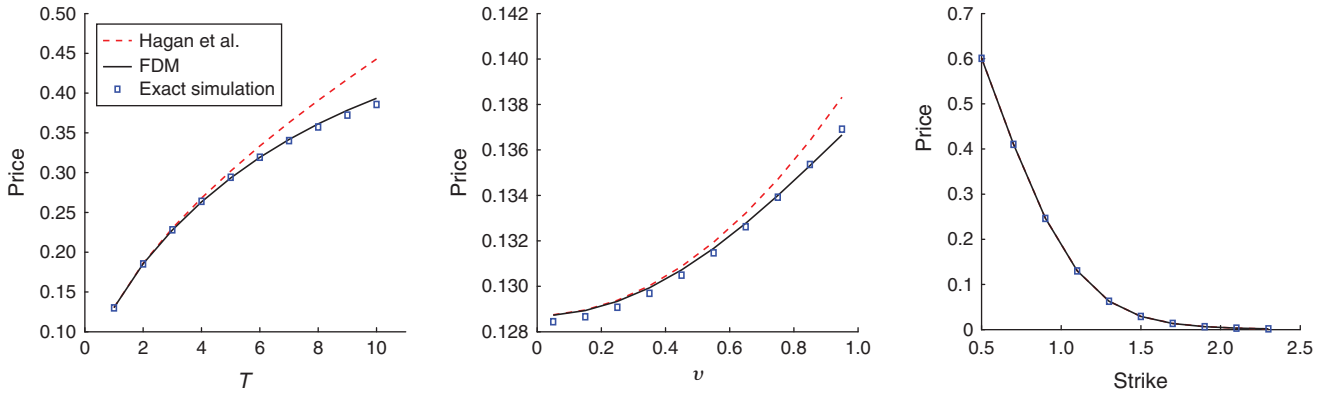

Notes. The sample size of the exact simulation is 100,000. We can see that our results are still accurate and coincide with those of the FDM.

Table 4. Parameter Settings in the Comparison between the Exact Simulation Method and the Euler Scheme

Case number	F_0	Range of K	T	ρ	β	ν	α_0
III.A	0.05	[0.01, 0.09]	1	0	0.30	0.60	0.40
III.B	0.05	[0.01, 0.09]	3	0	0.30	0.60	0.40
III.C	0.05	[0.01, 0.09]	5	0	0.30	0.60	0.40

Table 5 compares the two simulation methods through their RMS errors for the at-the-money European option prices in the three parameter settings in Table 4. Figure 3 illustrates how the RMS errors of the two simulation estimators converge. As suggested by Duffie and Glynn (1995), the number of time steps for the Euler scheme is simply set equal to the square root of the sample size. We can see that the RMS errors of the exact simulation estimators, which equal the standard errors, decay approximately at the rate $O(1/\sqrt{N})$ in all three cases; see Table 6 for the decay ratios as N increases and Figure 3 for its relationship with simulation times. In comparison, the RMS errors of the Euler estimators decay more slowly due to the biases.

4.2.2. Comparison with the Low-Bias Simulation Scheme of Chen et al. (2012). Recently, Chen et al. (2012) developed a low-bias simulation scheme for the SABR model. Note that their method is biased even when $\rho = 0$ because in an intermediate step, they employ the moment matching method to approximate the conditional distribution of $\int_0^T \alpha_s^2\, ds$ given α_0 and α_T. To reduce the bias incurred by the moment matching, they combine the moment matching with the discretization method. Therefore, like the Euler scheme, they need to increase the sample size and the number of time steps to improve the accuracy.

By contrast, our simulation scheme can exactly simulate $\int_0^T \alpha_s^2\, ds$, given α_0 and α_T, without biases. Table 7 compares the pricing results of the low-bias scheme, our exact simulation method, and the Euler scheme. The parameters are the same as in Case III.A, and the sample sizes are all 100,000. It can be seen that our exact simulation method is more efficient than the low-bias scheme and the Euler scheme, and requires less time to achieve the same accuracy.

Figure 3 shows the convergence of the RMS errors of the exact simulation, the Euler scheme, and the low-bias scheme in the parameter settings in Table 4. As suggested by Chen et al. (2012), the number of time steps of the low-bias scheme is chosen to be 1/50 of that for the Euler scheme. When $T = 1$, the slopes (computed by fitting with the least squares method) for the exact simulation, the Euler scheme, and the low-bias scheme are −0.499, −0.251, and −0.294, respectively. When $T = 3$, the slopes for the three methods are −0.495, −0.247, and −0.314, respectively. When $T = 5$, the slopes for the three methods are −0.493, −0.250, and −0.327, respectively. As one can see, the Euler and low-bias schemes converge more slowly than the exact simulation.

Table 5. Comparison of RMS Errors of at-the-Money European Call Option Prices between the Exact Simulation and the Euler Scheme in the Three Parameter Settings in Table 4

	Exact simulation		Euler scheme				
Sample size	RMS err.	Time (sec.)	Steps	Bias	Std. err.	RMS err.	Time (sec.)
Prices of European call options in Case III.A ($T = 1$)							
10,000	1.14E−03	10.0	100	3.07E−03	1.25E−03	3.32E−03	1.3
40,000	6.02E−04	44.6	200	1.83E−03	6.20E−04	1.93E−03	11.2
160,000	3.05E−04	179.0	400	1.14E−03	3.12E−04	1.18E−03	88.1
640,000	1.52E−04	715.0	800	7.01E−04	1.52E−04	7.17E−04	694.9
2,560,000	7.62E−05	2,482.7	1,600	3.77E−04	7.62E−05	3.85E−04	4,814.2
Prices of European call options in Case III.B ($T = 3$)							
10,000	1.82E−03	9.1	100	9.16E−03	2.41E−03	9.48E−03	1.4
40,000	1.02E−03	41.2	200	5.45E−03	1.09E−03	5.56E−03	10.7
160,000	5.23E−04	165.1	400	3.40E−03	6.15E−04	3.46E−03	84.0
640,000	2.63E−04	654.1	800	2.04E−03	2.74E−04	2.06E−03	675.5
2,560,000	1.34E−04	2,341.0	1,600	1.17E−03	1.33E−04	1.18E−03	4,707.7
Prices of European call options in Case III.C ($T = 5$)							
10,000	2.09E−03	11.7	100	1.44E−02	2.94E−03	1.47E−02	1.5
40,000	1.27E−03	42.2	200	8.51E−03	1.39E−03	8.62E−03	11.6
160,000	6.62E−04	168.9	400	5.32E−03	9.13E−04	5.39E−03	90.3
640,000	3.40E−04	676.1	800	3.15E−03	3.66E−04	3.17E−03	711.5
2,560,000	1.77E−04	2,419.8	1,600	1.87E−03	1.76E−04	1.88E−03	5,001.2

Figure 3. (Color online) Convergence of the RMS Errors of the Exact Simulation Method, the Euler Scheme, and the Low-Bias Scheme in the Three Parameter Settings in Table 4, Namely, Cases III.A–C

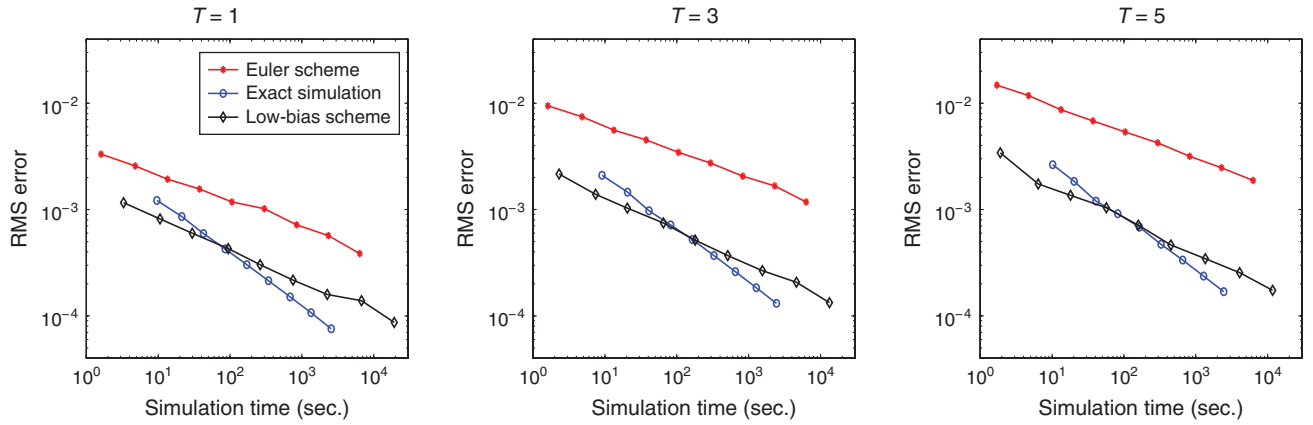

Notes. When $T = 1$, the slopes for the exact simulation method, the Euler scheme, and the low-bias scheme are −0.499, −0.251, and −0.294, respectively. When $T = 3$, the slopes for the three methods are −0.495, −0.247, and −0.314, respectively. When $T = 5$, the slopes for the three methods are −0.493, −0.250, and −0.327, respectively.

4.3. A PSE Simulation Method When $\rho \neq 0$

When $\beta < 1$ and $\rho \neq 0$, our simulation method becomes a semi-exact one because the samplings of α_T in Step 1 and $\int_0^T \alpha_s^2\, ds$ in Step 2 are exact, and only the simulation of F_T given α_T and $\int_0^T \alpha_s^2\, ds$ in Step 3 is based on an approximate conditional distribution. More precisely, we approximate α_u for any $u \in (0, T]$ in Equation (EC.10) by α_T to construct the approximate conditional distribution of F_T. See Section EC.2 in the e-companion for details. Therefore, intuitively, the approximation is expected to be good when T is small, and to become worse when T is big.

4.3.1. The Semi-Exact Simulation Scheme When T Is Small (e.g., $T \leqslant 1$).

To illustrate the performance of the semi-exact simulation scheme for small T, we use it to compute the European option prices for wide rages of parameters in the foreign exchange market, and compare them with those obtained via the FDM. Specifically, we take 3 values for each of the 3 model parameters, ν, β, and ρ, as 3 representative levels (low, middle, and high), and select 11 strike prices, including at, in and out of the money cases. The sample size of the simulation is 500,000. Numerical results suggest that the semi-exact simulation scheme is quite accurate when T is not big, e.g., $T \leqslant 1$, and becomes more accurate as T becomes smaller.

A summary of the results is given in Table 8. When $T = 1$, for most cases (92.59%), the prices obtained via the FDM fall into the narrow 95% CIs of our semi-exact simulation estimates. The FDM prices fall outside the CIs only in the extreme cases of $\nu = 0.9$ and $\rho = -0.9$ (i.e., the volatility process is highly volatile and highly correlated with the price process). However, in these extreme cases, the pricing errors are still quite small with the average relative error 1.33%. When $T = 0.5$, it turns out that the FDM prices in all 297 cases fall into the narrow CIs of our semi-exact simulation estimates.

4.3.2. A PSE Simulation Scheme When T Is Big.

When the maturity T becomes large, our semi-exact simulation method may no longer perform well. As one can see in Figure 4, as T rises, our semi-exact simulation method and the expansion formula of Hagan et al.

Table 6. How the RMS Errors (Namely, Standard Errors) of the Exact Simulation Estimators Converge to 0 in the Three Parameter Settings in Table 4

The convergence rate of the exact simulation estimator						
	Case III.A		Case III.B		Case III.C	
Sample size	RMS err.	Decay ratio	RMS err.	Decay ratio	RMS err.	Decay ratio
10,000	1.14E−03	N/A	1.82E−03	N/A	2.09E−03	N/A
40,000	6.02E−04	0.526	1.02E−03	0.561	1.27E−03	0.609
160,000	3.05E−04	0.506	5.23E−04	0.514	6.62E−04	0.521
640,000	1.52E−04	0.500	2.63E−04	0.503	3.40E−04	0.513
2,560,000	7.62E−05	0.500	1.34E−04	0.508	1.77E−04	0.519

Note. It can be seen that the convergence rate is approximately $O(1/\sqrt{N})$, and verifies that our exact simulation estimators are unbiased.

Table 7. Comparison of European Call Option Prices Produced by the Exact Simulation, the Low-Bias Scheme, and the Euler Scheme

		K/F_0						
	Δt	40%	80%	100%	120%	160%	200%	Time (sec.)
Euler	1/400	0.0472	0.0429	0.0409	0.0389	0.0352	0.0318	49.4
	1/800	0.0463	0.0420	0.0399	0.0380	0.0343	0.0309	99.1
	1/1,600	0.0453	0.0411	0.0391	0.0372	0.0336	0.0303	194.0
Low-bias	1/4	0.0461	0.0419	0.0399	0.0379	0.0343	0.0310	78.4
	1/8	0.0460	0.0418	0.0398	0.0378	0.0342	0.0308	175.8
Exact		0.0457	0.0416	0.0396	0.0377	0.0341	0.0308	98.3
True values		0.0456	0.0414	0.0394	0.0375	0.0339	0.0306	

Notes. Parameters are the same as in Case III.A. "True values" are obtained by the FDM.

Table 8. A Summary of the Numerical Results for Testing the Semi-Exact Simulation Scheme When $T = 1$ and 0.5

$T = 1$		$T = 0.5$	
The percentage of the cases where the FDM prices fall into the CIs	92.59%	The percentage of the cases where the FDM prices fall into the CIs	100.00%
The average width of CIs	1.07E−03	The average width of CIs	7.20E−04
The average absolute error	2.21E−04	The average absolute error	7.16E−05
The maximum absolute error	2.54E−03	The maximum absolute error	2.36E−04
The average relative error	0.19%	The average relative error	0.08%
The maximum relative error	3.17%	The maximum relative error	0.34%
The average relative error for all "outside the CIs" cases	1.33%		

Notes. The parameters $\nu \in \{0.1, 0.5, 0.9\}$, $\beta \in \{0.1, 0.5, 0.9\}$, $\rho \in \{-0.1, -0.5, -0.9\}$, $K \in \{1.00, 1.02, \ldots, 1.20\}$, $F_0 = 1.1$, and $\alpha_0 = 0.3$. "CI" denotes "confidence interval." Absolute and relative errors are between the FDM prices and our simulation estimates. We can see that the semi-exact simulation scheme is quite accurate when $T \leqslant 1$, and becomes more accurate as T becomes smaller.

Figure 4. (Color online) European Option Prices Obtained by Our Semi-Exact Simulation Method, Our PSE Simulation Method (with $\Delta = 1$ and 0.5), the Expansion Formula of Hagan et al. (2002) and the FDM, When the Maturity T Increases

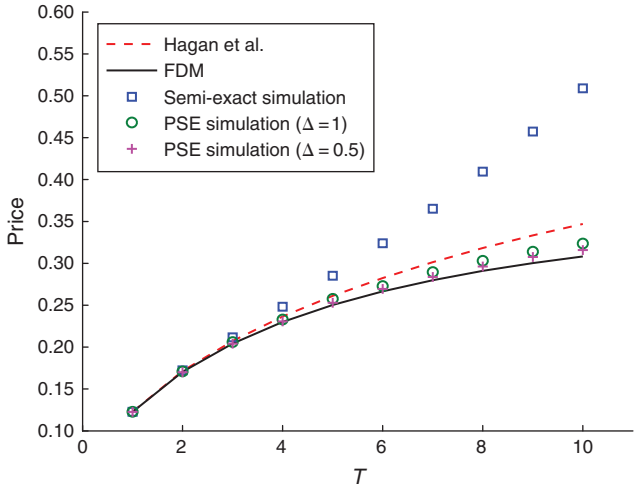

Notes. The parameters are $F_0 = 1.10$, $\rho = -0.8$, $\beta = 0.4$, $\nu = 0.5$, and $\alpha_0 = 0.3$. The sample size is 100,000. We can see that the PSE simulation becomes more and more accurate when Δ decreases from T (i.e., the semi-exact simulation), to 1, and then to 0.5.

(2002) cause increasing upward biases. Moreover, the biases of our method are even bigger than those of the expansion formula of Hagan et al. (2002).

To reduce the biases, we intend to propose a PSE simulation method. The idea is to first divide the time horizon $[0, T]$ into several short pieces with the same piece length Δ, and then to simulate F_T by applying the semi-exact simulation method piece by piece. Since our semi-exact simulation scheme is quite accurate for short time intervals, this PSE simulation algorithm is expected to perform well when the piece length Δ is sufficiently small. It is worth pointing out that Δ needs not to be very small (i.e., the number of pieces needs not to be big). In fact, $\Delta = 1$ or 0.5 has been short enough to achieve a satisfactory accuracy. As demonstrated in Figure 4, the PSE simulation scheme becomes increasingly accurate when Δ decreases from T (i.e., the semi-exact simulation), to 1, and then to 0.5. Indeed, when $\Delta = 1$, our method has outperformed the expansion formula of Hagan et al. (2002) and has been quite accurate when $T \leqslant 6$. When $\Delta = 0.5$, our method is quite accurate even in the extreme case $T = 10$. In summary, we take the following steps to implement our PSE simulation scheme.

The PSE Simulation Method

Step 1. Divide the time horizon into a small number of short pieces with the piece length Δ (typically $\Delta = 1$ or 0.5 is short enough in practice);

Step 2. Apply the semi-exact simulation method developed in Section 3 piece by piece until we obtain a sample of F_T.

Remark 4.1. It is worth noting that the PSE simulation method differs greatly from the common discretization in simulation. First, the former only needs to divide the time horizon into a very small number of pieces, e.g., five pieces when $T = 5$. However, the latter always requires a much larger number (e.g., hundreds or thousands) of discretization time steps. Second, in each time step, the common discretization simply uses the normal distribution to approximate the targeted one, while in each piece, our PSE simulation applies a sophisticated semi-exact scheme to accomplish the simulation.

4.4. Comparison between the PSE Simulation Method and the Euler Scheme as Well as the Low-Bias Scheme

This section is devoted to comparing our PSE simulation method with the Euler scheme and the low-bias scheme. First, we would like to compare the biases of these three simulation methods. As mentioned in Remark 4.1, for the same maturity, our PSE simulation method requires a much smaller number of time pieces (or time steps) to reduce the bias to an acceptable level than does the common discretization method such as the Euler scheme (a similar phenomenon is also observed in the comparison between the low-bias scheme and the Euler scheme in Chen et al. 2012). To further support this, we conduct a numerical comparison of the biases of the European option prices obtained from the PSE simulation method, the Euler scheme and the low-bias scheme with respect to the number of time steps, and the results are shown in Figure 5. As we can see, on the one hand, the bias of the PSE simulation method decays faster than the Euler scheme. Indeed, if we use N_{time} to denote the number of time steps, Figure 5 suggests that the bias of the PSE simulation method is approximately proportional to $(N_{\text{time}})^{-2}$, while the bias of the Euler scheme is approximately proportional to $(N_{\text{time}})^{-1}$. In contrast, although the bias of the low-bias scheme appears to have a similar decay rate as that of our PSE simulation method, the magnitude of its bias is bigger than ours for the same number of time steps. Intuitively, this is because the low-bias scheme incurs an extra bias (caused by the moment matching) in addition to the same bias as that for our PSE simulation method.

Figure 5. (Color online) Decay of the Biases of the European Option Prices Obtained from the PSE Simulation Method, the Euler Scheme and the Low-Bias Scheme with Respect to the Number of Time Steps

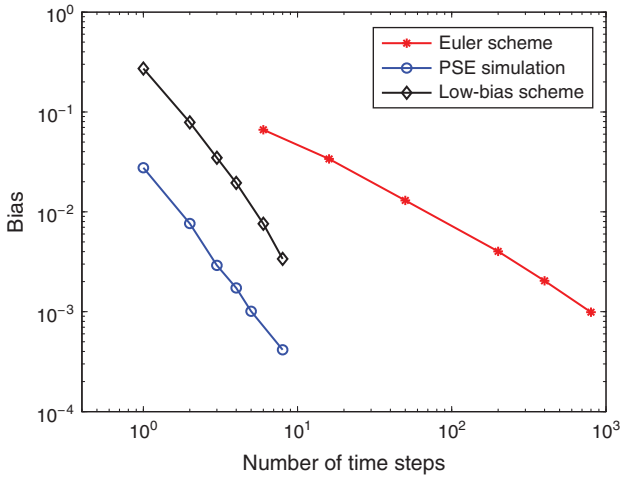

Notes. The slopes for the PSE method, the Euler scheme, and the low-bias scheme are about −2.0, −0.9, and −2.0, respectively. The parameters are $F_0 = 1.10$, $\rho = -0.3$, $\beta = 0.3$, $\nu = 0.8$, $\alpha_0 = 0.4$, and $T = 4$. The benchmark prices are obtained by the FDM. The sample sizes for the bias estimation of these three methods are 10,240,000.

Next, we would like to compare the RMS errors of these three simulation methods. Like the Euler scheme, when applying our PSE simulation method and the low-bias scheme to option pricing under a limited computational budget, we also need to make a trade-off between increasing the number of time steps N_{time} to reduce the bias and increasing the number of simulation samples N to reduce the statistical error. Note that as shown in Figure 5, the biases of our PSE method and the low-bias scheme decay faster than that of the Euler scheme. Intuitively, this suggests that it might be possible to enhance the overall efficiency of these two methods by spending less computational effort on the bias reduction. Indeed, Duffie and Glynn (1995) show that for a simulation scheme designed for option pricing, if the bias decays at the order of $(N_{\text{time}})^{-p}$, then asymptotically it is optimal to increase N_{time} proportional to $N^{1/(2p)}$. Accordingly, for the Euler scheme, Duffie and Glynn (1995) suggest choosing the number of time steps proportional to the square root of the sample size (also see, e.g., Broadie and Kaya 2006). Likewise, for our PSE simulation method and the low-bias scheme, we would like to have N_{time} increasing at the order of $N^{1/4}$. This implies less computational effort on the bias reduction compared to the Euler scheme, which coincides with our intuition. As mentioned in Broadie and Kaya (2006), the optimal constant of proportionality is usually not easy to determine. However, this does not affect the asymptotically optimal decay rate of the RMS error. According to the suggestion of Broadie and Kaya (2006), for a given T and a specific

Table 9. Comparison of the RMS Errors of the European Option Prices Obtained from the PSE Simulation Method, the Euler Scheme, and the Low-Bias Scheme in the Same Parameter Setting as in Figure 5

	PSE simulation			Euler scheme			Low-bias scheme		
Sample size	Steps	RMS err.	Time (sec.)	Steps	RMS err.	Time (sec.)	Steps	RMS err.	Time (sec.)
10,000	1	2.85E–02	14.3	400	7.62E–03	7.1	2	7.90E–02	9.0
40,000	2	8.46E–03	116.0	800	3.54E–03	55.2	3	3.50E–02	61.2
160,000	3	3.44E–03	655.6	1,600	1.77E–03	441.4	4	1.96E–02	349.8
640,000	4	1.99E–03	3,865.5	3,200	1.09E–03	3,500.7	6	7.66E–03	2,266.4
2,560,000	6	5.03E–04	22,597.7	6,400	5.92E–04	27,002.6	8	3.41E–03	12,631.9

sample size N, the number of time steps N_{time} of our PSE simulation method should be selected such that the discretization bias is close to the statistical error. Then, if we expect to further improve the accuracy, we can increase N_{time} proportional to $N^{1/4}$.

Table 9 compares the RMS errors of the three simulation methods for the at-the-money European option prices in the same parameter setting as that in Figure 5. The numbers of time steps are chosen based on the discussion above. Figure 6 illustrates how the RMS errors of the three simulation estimators converge. In comparison, the RMS error of our PSE simulation method decays faster than that of the Euler scheme. Moreover, to achieve the same level of the RMS error, it takes less time to use our PSE simulation method than it does to use the low-bias scheme.

4.5. An Example of Pricing Exotic Options: Forward Starting Options

Since analytical solutions or analytical approximations usually do not exist for exotic option prices under the SABR model, the Monte Carlo simulation approach becomes especially useful. Here, we give an example to price one type of exotic options—forward starting options—through our exact (or semi-exact) simulation method. A forward starting option with maturity T_2 is the same as a European call option with a floating strike kS_{T_1}, where k is a positive constant and S_{T_1} is the underlying asset price at an intermediate time $T_1 \in (0, T_2)$. Therefore, its price is given by $\mathbb{E}[e^{-rT_2}(S_{T_2} - kS_{T_1})^+]$. The simulation method applies in a straightforward way. Table 10 gives the RMS errors of the forward starting option prices obtained by the exact simulation method and the Euler scheme in the same parameter settings as in Table 4, while Figure 7 illustrates the convergence of the RMS errors for these two methods. We can see that the RMS errors of the exact simulation estimators decay faster than those of the Euler estimators.

Figure 6. (Color online) Convergence of the RMS Errors of the PSE Simulation Method, the Euler Scheme, and the Low-Bias Scheme in the Same Parameter Setting as in Figure 5

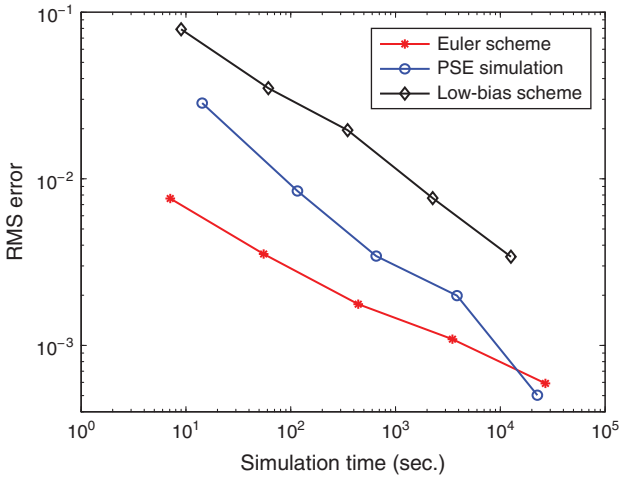

5. A Conditional Simulation Method for European Option Pricing

This section proposes a conditional simulation method for the pricing of European options under the SABR model, which can substantially reduce the variance (by, e.g., 99%) of the plain simulation estimators developed in Section 3. This method is based on the following theoretical result.

Theorem 5.1. *The European call option price* $\mathbb{E}[(F_T - K)^+]$ *under the SABR model with the absorbing boundary* 0 *has the following formulas*:

(i) *When* $\beta = 1$,

$$\mathbb{E}[(F_T - K)^+] = F_0 \mathbb{E}[e^{-(\rho^2/2)\int_0^T \alpha_s^2\, ds + (\rho/\nu)(\alpha_T - \alpha_0)} \Psi(d_1)] - K\mathbb{E}[\Psi(d_2)], \tag{20}$$

where $\Psi(\cdot)$ *denotes the standard normal cdf,*

$$d_1 = \frac{(1/2 - \rho^2)\int_0^T \alpha_s^2\, ds + (\rho/\nu)(\alpha_T - \alpha_0) - \log(K/F_0)}{\sqrt{(1-\rho^2)\int_0^T \alpha_s^2\, ds}}$$

and

$$d_2 = \frac{-(1/2)\int_0^T \alpha_s^2\, ds + (\rho/\nu)(\alpha_T - \alpha_0) - \log(K/F_0)}{\sqrt{(1-\rho^2)\int_0^T \alpha_s^2\, ds}}.$$

(ii) *When* $\beta \in [0, 1)$ *and* $\rho = 0$,

$$\mathbb{E}[(F_T - K)^+] = \mathbb{E}[C_0(1)^{-1/(2(1-\beta))} A_0^{(1+\gamma_0)/2} (1 - Q(C_0(K); 3 + \gamma_0, A_0))] - K\mathbb{E}[Q(A_0; 1 + \gamma_0, C_0(K))], \tag{21}$$

Table 10. Comparison of the RMS Errors of the Forward Starting Option Prices Obtained from the Exact Simulation Method and the Euler Scheme

	Exact simulation		Euler scheme				
Sample size	RMS err.	Time (sec.)	Steps	Bias	Std. err.	RMS err.	Time (sec.)
10,000	1.06E–03	12.5	100	1.90E–03	9.55E–04	2.13E–03	2.1
40,000	4.93E–04	49.3	200	1.16E–03	6.07E–04	1.31E–03	16.7
160,000	2.52E–04	198.3	400	7.32E–04	2.80E–04	7.83E–04	130.0
640,000	1.29E–04	797.8	800	4.53E–04	1.29E–04	4.71E–04	1,027.3
2,560,000	6.37E–05	3,198.5	1,600	2.60E–04	6.44E–05	2.68E–04	8,185.6

Note. The parameters used are the same as in Table 4 except that $T_1 = 1$, $T_2 = 2$, and $k = 1$.

where $\gamma_0 = \beta/(1-\beta)$, and A_0 and $C_0(u)$ are the same as in Proposition 2.1.

(iii) *When $\beta \in [0,1)$ and $\rho \neq 0$,*

$$\mathbb{E}[(F_T-K)^+] \approx \mathbb{E}\left[C(1)^{-1/(2(1-\beta))}A^{(1+\gamma)/2}\Phi^+\left(-\frac{\rho^2\gamma}{2}, C(K); 3+\gamma, A\right)\right] - K\mathbb{E}[Q(A; 1+\gamma, C(K))],$$

where $\gamma = \beta/((1-\beta)(1-\rho^2))$, A and $C(u)$ are the same as in Proposition 2.1, and

$$\Phi^+(p,k;\delta,\alpha) := 2^p \sum_{n=0}^{+\infty} e^{-\alpha/2}\left(\frac{\alpha}{2}\right)^n \frac{\Gamma(\delta/2+p+n, k/2)}{n!\,\Gamma(\delta/2+n)} \quad (22)$$

with $\Gamma(a,x)$ denoting the complementary incomplete gamma function.

According to Theorem 5.1, we propose the following conditional simulation method for European option pricing, which is exact when $\beta = 1$, or $\beta \in [0,1)$ and $\rho = 0$, and is semi-exact when $\beta \in [0,1)$ and $\rho \neq 0$.

Figure 7. (Color online) Convergence of the RMS Errors of the Exact Simulation Method and the Euler Scheme for Pricing of the Forward Starting Options

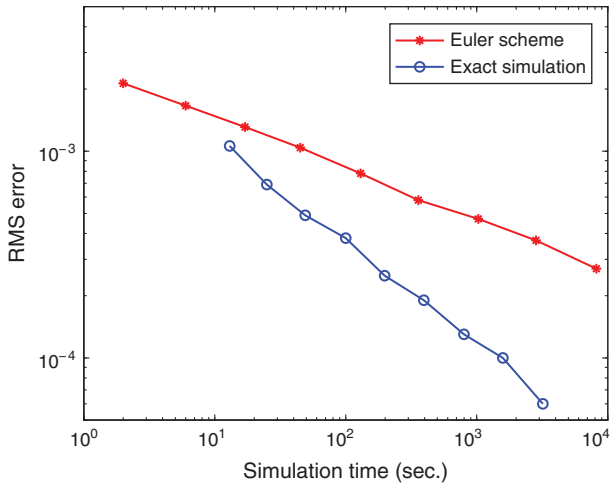

Note. The slopes for the exact simulation method and the Euler scheme are about −0.492 and −0.250, respectively.

The Conditional Simulation Method for European Option Pricing

Step 1. Exactly generate a sample pair α_T and $\int_0^T \alpha_s^2\,ds$, given α_0, using the exact simulation method developed in Section 3;

Step 2. A sample of $\mathbb{E}[(F_T-K)^+ \mid \alpha_0, \alpha_T, \int_0^T \alpha_s^2\,ds, F_0]$ is given by $F_0\, e^{-(\rho^2/2)\int_0^T \alpha_s^2\,ds + (\rho/\nu)(\alpha_T-\alpha_0)}\,\Psi(d_1) - K\Psi(d_2)$ if $\beta = 1$, by $C_0(1)^{1/(2(\beta-1))}A_0^{(1+\gamma_0)/2}(1 - Q_{\chi'^2}(C_0(K); 3+\gamma_0, A_0)) - K\,Q_{\chi'^2}(A_0; 1+\gamma_0, C_0(K))$ if $\beta \in [0,1)$ and $\rho = 0$, and approximately by $C(1)^{1/(2(\beta-1))}A^{(1+\gamma)/2}\,\Phi^+(-\rho^2\gamma/2, C(K); 3+\gamma, A) - KQ_{\chi'^2}(A; 1+\gamma, C(K))$ if $\beta \in [0,1)$ and $\rho \neq 0$;

Step 3. Repeat Steps 1 and 2 for N times, and the sample average gives an estimate for the European option price $\mathbb{E}[(F_T-K)^+]$.

Remark 5.1. Broadie and Kaya (2006) apply the conditional simulation method for European option pricing under the Heston model. They show that the conditional European option price, given the volatility process, can be expressed as a Black-Scholes formula, thanks to the conditional normal distribution of the asset price. As for the SABR model, when $\beta = 1$, the forward price F_T is also normally distributed conditional on α_T and $\int_0^T \alpha_s^2\,ds$. Hence, a similar conditional Black-Scholes formula is also available. However, when $\beta < 1$, the conditional distribution of F_T becomes much more complicated. We manage to derive an explicit formula (exact if $\rho = 0$ and approximate if $\rho \neq 0$) for the conditional European option price given α_T and $\int_0^T \alpha_s^2\,ds$, using the results about the truncated moments of noncentral chi-squared random variables in Carr and Linetsky (2006).

The conditional simulation scheme enables us to eliminate part of the variance caused by the price process. Specifically, it removes the randomness of $\{W_t^{(1)}: 0 \leqslant t \leqslant T\}$, hence achieving a significant variance reduction compared to the plain simulation method in Section 3. To illustrate, Table 11 provides a comparison of the numerical prices produced by the conditional simulation scheme and the plain exact simulation. It

Table 11. Variance Reduction of the Conditional Simulation Method for European Option Pricing Compared to the Plain Exact Simulation

	Exact simulation			Conditional exact simulation			
Sample size	Price	Std. err.	Time (sec.)	Price	Std. err.	Time (sec.)	Var. red. (%)
			Prices of European call options in Case III.A ($T = 1$)				
10,000	0.03763	1.14E–03	10.0	0.03945	4.08E–05	8.6	99.87
40,000	0.03950	6.02E–04	44.6	0.03944	2.04E–05	34.5	99.89
160,000	0.03965	3.05E–04	179.0	0.03942	1.03E–05	155.5	99.89
640,000	0.03960	1.52E–04	715.0	0.03941	5.14E–06	615.7	99.89
2,560,000	0.03951	7.62E–05	2,482.7	0.03942	2.57E–06	2,288.3	99.89
			Prices of European call options in Case III.B ($T = 3$)				
10,000	0.04075	1.82E–03	9.1	0.04368	3.95E–05	8.8	99.95
40,000	0.04390	1.02E–03	41.2	0.04368	1.98E–05	39.3	99.96
160,000	0.04394	5.23E–04	165.1	0.04364	9.96E–06	151.0	99.96
640,000	0.04377	2.63E–04	654.1	0.04364	4.99E–06	604.2	99.96
2,560,000	0.04380	1.34E–04	2,341.0	0.04364	2.50E–06	2,243.0	99.97
			Prices of European call options in Case III.C ($T = 5$)				
10,000	0.04055	2.09E–03	11.7	0.04473	3.86E–05	9.6	99.97
40,000	0.04460	1.27E–03	42.2	0.04472	1.93E–05	38.5	99.98
160,000	0.04485	6.62E–04	168.9	0.04469	9.76E–06	154.4	99.98
640,000	0.04469	3.40E–04	676.1	0.04469	4.89E–06	616.8	99.98
2,560,000	0.04481	1.77E–04	2,419.8	0.04469	2.45E–06	2,292.4	99.98

Notes. The parameters are the same as in Table 4 except that $K = 0.05$. The last column "Var. red." denotes the ratio of variance reduction defined as 1 – Var[Conditional simulation]/Var[Exact simulation]. We can see that the conditional simulation method with the same sample size can reduce the variance dramatically by, e.g., more than 99%.

can be seen that with the same sample size, the conditional method reduces the variances of the plain exact simulation estimators by more than 99%. We also compare the RMS errors of the conditional simulation with those of the plain exact simulation and the Euler scheme, and find that the conditional simulation method is much more efficient (see Figure 8).

6. An Extension: Exact Simulation of the SABR Model with a Reflecting Boundary

Under the SABR model with zero specified as a reflecting boundary, no analytical asymptotic expansions exist for the density function and distribution function of the forward price F_T. Moreover, analytical solutions or approximations for the European option price and the implied volatility are also unavailable. Consequently, Monte Carlo simulation becomes practically useful. The exact and semi-exact simulation methods proposed earlier for the SABR model with an absorbing boundary can be extended to the SABR model with a reflecting boundary based on the following result. See Islah (2009) or Section EC.2 in the e-companion for the proof.

Proposition 6.1 (Islah 2009). *Assume that* $T > 0, \beta \in [0, 1)$, 0 *is a reflecting boundary, and* α_0, α_T, $\int_0^T \alpha_u^2\, du$, *and* F_0 *are given. Define* $\delta_0 = (1-2\beta)/(1-\beta)$ *and* $\delta = 1-\beta/((1-\beta)\cdot(1-\rho^2))$.

Figure 8. (Color online) Convergence of the RMS Errors of the Conditional Simulation Method, the Plain Exact Simulation and the Euler Scheme in the Three Parameter Settings in Table 4, Namely, Cases III.A–C

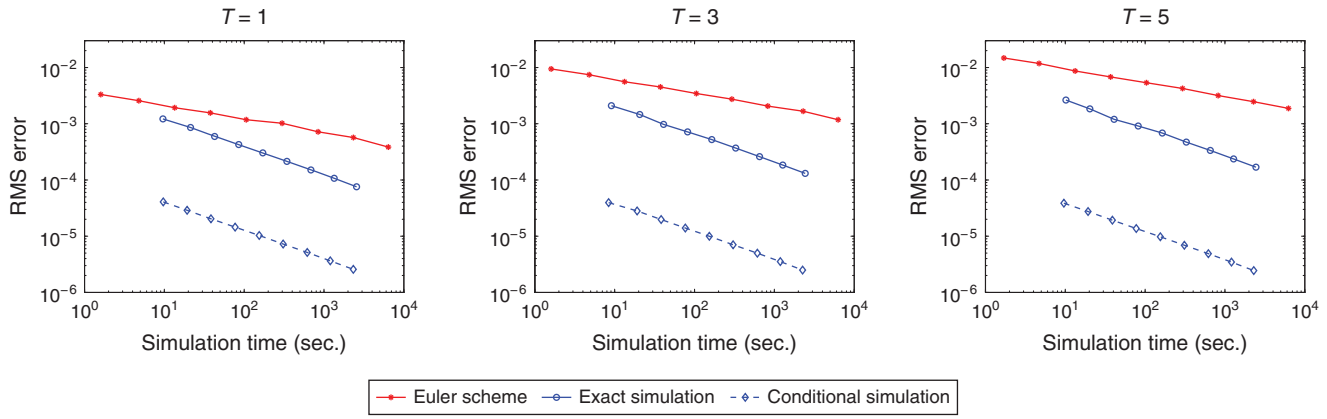

(i) *When* $\rho = 0$ *and* $\delta_0 \geqslant 0$,

$$\mathbb{P}\left(F_T \leqslant u \,\middle|\, F_0, \alpha_0, \alpha_T, \int_0^T \alpha_s^2\, ds\right) = Q_{\chi'^2}(C_0(u); \delta_0, A_0), \tag{23}$$

where $C_0(u)$ *and* A_0 *are the same as in Proposition 2.1. Indeed,* F_T *conditional on* α_0, α_T, $\int_0^T \alpha_s^2\, ds$, *and* F_0 *has the same distribution as the power of a scaled noncentral chi-squared random variable*

$$\left[\left((1-\beta)^2 \int_0^T \alpha_s^2\, ds\right) \cdot \chi'^2\left(\delta_0; A\right)\right]^{1/(2(1-\beta))}. \tag{24}$$

(ii) *When* $\rho \neq 0$ *and* $\delta \geqslant 0$,

$$\mathbb{P}\left(F_T \leqslant u \,\middle|\, F_0, \alpha_0, \alpha_T, \int_0^T \alpha_s^2\, ds\right) \approx Q_{\chi'^2}(C(u); \delta, A), \tag{25}$$

where $C(u)$ *and* A *are the same as in Proposition 2.1.*

Accordingly, we can take the same three steps as in Section 3 to exactly or semi-exactly sample from the distribution of F_T, except that in Step 3, we need to deal with a different random variable given in Proposition 6.1. Since in this case, the conditional cdf of F_T can also be expressed in terms of the noncentral chi-squared cdf, the inverse transform method used in Step 3 in the case of absorbing boundary still applies. Alternatively, (24) implies that F_T conditional on α_0, α_T, $\int_0^T \alpha_u^2\, du$, and F_0 can be simulated by first generating a sample for the noncentral chi-squared random variable $\chi'^2(\delta_0; A)$, and then substituting it into (24). We refer to Glasserman (2004) for the details of simulating a noncentral chi-squared random variable. Numerical experiments indicate that the two methods perform similarly well.

As an illustrative example, Figure 9 depicts the marginal density of F_T under the SABR model with a reflecting boundary using the exact simulation and the Euler scheme with different numbers of time steps (the right panel is the magnified part near zero of the left one). In the Euler scheme, once a negative value is generated in an intermediate step, we use its absolute value instead to guarantee that the generated value stays within the domain as well as to approximate the reflecting behavior. We can see that the densities generated by the two simulation methods are quite different in the neighborhood of 0. The density generated by the exact simulation tends to increase as F_T goes to 0, while the densities generated by the Euler scheme tend to decrease. However, Proposition 6.2 shows that when $\delta_0 \in (0, 1)$, the marginal density of F_T goes to $+\infty$ as u goes to 0. This implies that the exact simulation method performs better than the Euler scheme. Moreover, we can see that as the number of time steps increases, the density generated by the Euler scheme approaches to the density generated by the exact simulation.

Proposition 6.2. *Consider a SABR model with* $\rho = 0$, $\delta_0 \in (0, 1)$, *and a reflecting boundary* 0. *Denote by* $p_{F_T}(u)$ *the marginal density of* F_T. *Then,* $\lim_{u \to 0+} p_{F_T}(u) = +\infty$.

Remark 6.1. As documented in Goldstein and Keirstead (1997), in the period of 1991–1995, the short-term interest rate in Japan declined to a very low level (below 0.4%) and still remained highly volatile. To address this issue, it was proposed to describe the interest rates using models with reflecting boundaries (see, e.g., Goldstein and Keirstead 1997), and these studies were supported by much empirical evidence (see, e.g., Aït Sahalia 1996, Stanton 1997). Since an important application of the SABR model is in the interest rate markets (see, e.g., the recent book by Rebonato et al. 2011), it might be interesting to consider a SABR model

Figure 9. (Color online) Marginal Densities of F_T Under the SABR Model with a Reflecting Boundary

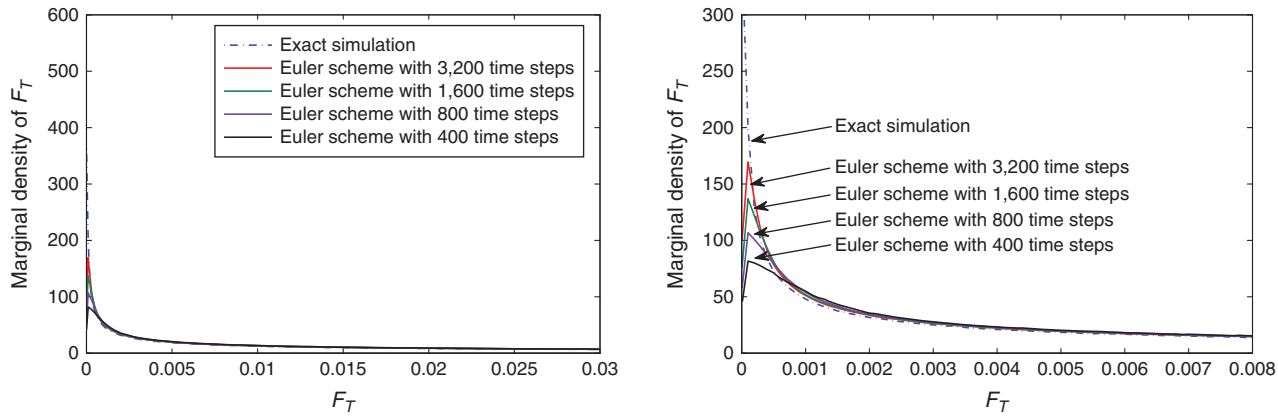

Notes. The parameters are $\nu = 0.04$, $\beta = 0.3$, $\rho = 0$, $F_0 = 0.05$, $\alpha_0 = 0.2$, and $T = 1$. The sample size is 10,240,000. We can see that the densities generated by the exact simulation and the Euler scheme are quite different near zero. Moreover, as the number of time steps increases, the density generated by the Euler scheme approaches to the density generated by the exact simulation. The right panel is the magnified part near zero of the left one.

with reflecting boundary for the modeling of the interest rates in such markets as Japan. Since the interest rates in these markets could be very small (close to zero) and still have great volatility, an accurate simulation in the neighborhood of zero of the SABR model with the reflecting boundary becomes important.

7. Concluding Remarks

In this paper, we explore the possibility of exact simulation for the SABR model, one of the most popular stochastic volatility models in finance. In two special but practically interesting cases, i.e., when the elasticity $\beta = 1$, or $\beta < 1$ and the price and volatility processes are instantaneously uncorrelated ($\rho = 0$), we develop an exact simulation scheme for the forward price and its volatility by proposing two novel simulation methods to circumvent two involved difficulties. When $\beta < 1$ and $\rho \neq 0$, our simulation method becomes a semi-exact one, which turns out to be still quite accurate when the time horizon is not long, e.g., no greater than one year. When the time horizon becomes longer, the semi-exact simulation incurs increasing biases. A PSE simulation scheme is developed that reduces the biases substantially. Finally, we propose a conditional simulation method for European option pricing under the SABR model. Numerical results suggest that it can reduce the variance of the plain simulation dramatically, e.g., by more than 99%.

Acknowledgments

The authors are grateful to the editor, the area editor, the associate editor, and two anonymous referees for their helpful comments.

References

Abate J, Whitt W (1992) The Fourier-series method for inverting transforms of probability distributions. *Queueing Systems* 10(1):5–88.

Aït Sahalia Y (1996) Testing continuous-time models of the spot interest rate. *Rev. Finance Stud.* 9(2):385–426.

Andersen L, Andreasen J (2000) Volatility skews and extension of the LIBOR market model. *Appl. Math. Finance* 7(1):1–32.

Barndorff-Nielsen OE, Shephard N (2001) Non-Gaussian Ornstein-Uhlenbeck-based models and some of their uses in financial economics. *J. Roy. Statist. Soc., Ser. B.* 63(2):167–241.

Barrieu P, Rouault A, Yor M (2004) A study of the Hartman-Watson distribution motivated by numerical problems related to the pricing of Asian options. *J. Appl. Probab.* 41(4):1049–1058.

Benaim S, Friz P (2009) Regular variation and smile asymptotics. *Math. Finance* 19(1):1–12.

Berestycki H, Busca J, Florent I (2004) Computing the implied volatility in stochastic volatility models. *Comm. Pure Appl. Math.* 57(10):1352–1373.

Beskos A, Roberts G (2005) Exact simulation of diffusions. *Ann. Appl. Probab.* 15(4):2422–2444.

Borodin AN, Salminen P (2002) *Handbook of Brownian Motion—Facts and Formulae* (Springer, Berlin).

Boyle P, Potapchik A (2006) Application of high-precision computing for pricing arithmetic Asian options. *Internat. Symp. Symbolic Algebraic Comput.*, 39–46.

Broadie M, Kaya O (2004) Exact simulation of option Greeks under stochastic volatility and jump diffusion models. Ingalls RG, Rossetti MD, Smith JS, Peters BA, eds. *Proc. 2004 Winter Simulation Conf.* (IEEE, Piscataway, NJ), 1607–1615.

Broadie M, Kaya O (2006) Exact simulation of stochastic volatility and other affine jump diffusion processes. *Oper. Res.* 54(2): 217–231.

Carr P, Linetsky V (2006) A jump to default extended CEV model: An application of Bessel processes. *Finance and Stochastics* 10(3): 303–330.

Chen N, Huang Z (2013) Localization and exact simulation of Brownian motion driven stochastic differential equations. *Math. Oper. Res.* 38(3):591–616.

Chen B, Oosterlee CW, van der Weide H (2012) A low-bias simulation scheme for the SABR stochastic volatility model. *Internat. J. Theor. Appl. Finance* 15(2):125–161.

Davydov D, Linetsky V (2001) Pricing and hedging path-dependent options under the CEV process. *Management Sci.* 47(7):949–965.

Derman E, Kani I (1994) Riding on a smile. *Risk* 7:32–39.

Derman E, Kani I (1998) Stochastic implied trees: Arbitrage pricing with stochastic term and strike structure of volatility. *Internat. J. Theor. Appl. Finance* 1:61–110.

Ding CG (1992) Algorithm AS 275: Computing the non-central χ^2 distribution function. *Appl. Statist.* 41(2):478–482.

Doust P (2012) No-arbitrage SABR. *J. Comput. Finance* 15(3):3–31.

Duffie D, Glynn P (1995) Efficient Monte Carlo estimation of security prices. *Ann. Appl. Probab.* 5(4):897–905.

Duffie D, Pan J, Singleton K (2000) Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* 68(6):1343–1376.

Duffy DJ (2006) *Finite Difference Methods in Financial Engineering: A Partial Differential Equation Approach* (John Wiley & Sons, Chichester, UK).

Dupire B (1994) Pricing with a smile. *Risk* 7:18–20.

Dupire B (1997) Pricing and hedging with smiles. Dempster MAH, Pliska SR, eds. *Mathematics of Derivative Securities* (Cambridge University Press, Cambridge, UK), 103–111.

Dyrting S (2004) Evaluating the noncentral chi-squared distribution for the Cox-Ingersoll-Ross process. *Comput. Econom.* 24(1):35–50.

Fouque J-P, Papanicolaou G, Sircar R (2000) *Derivatives in Financial Markets with Stochastic Volatility* (Cambridge University Press, Cambridge, UK).

Giesecke K, Smelov D (2013) Exact sampling of jump-diffusions. *Oper. Res.* 61(4):894–907.

Giesecke K, Kakavand H, Mousavi M (2011) Exact simulation of point processes with stochastic intensities. *Oper. Res.* 59(5): 1233–1245.

Glasserman P (2004) *Monte Carlo Methods in Financial Engineering* (Springer, New York).

Goldstein R, Keirstead W (1997) On the term structure of interest rates in the presence of reflecting and absorbing boundaries. Preprint.

Hagan PS, Kumar D, Lesniewski AS, Woodward DE (2002) Managing smile risk. *Wilmott Magazine* 84–108.

Henry-Labordère P (2009) *Analysis, Geometry and Modeling in Finance* (CRC Press, Boca Raton, FL).

Heston S (1993) A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Rev. Financial Stud.* 6(2):153–187.

Hull J, White A (1987) The pricing of options on assets with stochastic volatilities. *J. Finance* 42(2):281–300.

Hull J, White A (1988) An analysis of the bias in option pricing caused by a stochastic volatility. *Adv. Futures Options Res.* 3:29–61.

Islah O (2009) Solving SABR in exact form and unifying it with LIBOR market model. Working paper.

Johnson H, Shanno D (1987) Option pricing when the variance changes randomly. *J. Financial Quant. Anal.* 22(4):143–151.

Karatzas I, Shreve S (1992) *Brownian Motion and Stochastic Calculus* (Springer, Berlin).

Larguinho M, Dias JC, Braumann CA (2011) Speed and accuracy comparison of noncentral chi-square distribution methods for option pricing and hedging under the CEV Model. *Proc. 18th Internat. Conf. Forecasting Financial Markets: Advances for Exchange Rates, Interest Rates and Asset Management* (CD-ROM).

Matsumoto H, Yor M (2005) Exponential functionals of Brownian motion I: Probability laws at fixed time. *Probab. Surveys* 2: 312–347.

Obłój J (2008) Fine-tune your smile: Correction to Hagan et al. *Wilmott Magazine* 35:102–104.

Penev S, Raykov T (2000) A Wiener germ approximation of the noncentral chi-square distibution and of its quantiles. *Comput. Statist.* 15(2):219–228.

Rebonato R, McKay K, White R (2011) *The SABR/LIBOR Market Model: Pricing, Calibration and Hedging for Complex Interest-Rate Derivatives* (John Wiley & Sons, Chichester, UK).

Ren Y, Madan D, Qian MQ (2007) Calibrating and pricing with embedded local volatility models. *Risk* 20:138–143.

Sankaran M (1963) Approximations to the non-central chi-square distribution. *Biometrika* 50(1/2):199–204.

Schroder M (1989) Computing the constant elasticity of variance option pricing fromula. *J. Finance* 44(1):211–219.

Scott L (1987) Option pricing when the variance changes randomly: Theory, estimation, and an application. *J. Financial Quant. Anal.* 22(4):419–438.

Stanton R (1997) A nonparametric model of term structure dynamics and the market price of interest rate risk. *J. Finance* 52(5): 1973–2002.

Stein E, Stein J (1991) Stock price distributions with stochastic volatility: An analytic approach. *Rev. Financial Stud.* 4(4):727–752.

Wiggins J (1987) Option values under stochastic volatility: Theory and empirical estimates. *J. Financial Econom.* 19(2):351–372.

Wu Q (2012) Series expansion of the SABR joint density. *Math. Finance* 22(2):310–345.

Yor M (1980) Loi de l'indice du lacet Brownien, et distribution de Hartman-Watson. *Z. Wahr. Verw. Gebiete* 53(1):71–95.

Yor M (1992) On some exponential functionals of Brownian motion. *J. Appl. Probab.* 24(3):509–531.

Ning Cai is an associate professor in the Department of Industrial Engineering and Logistics Management at the Hong Kong University of Science and Technology. His research interests include financial engineering, risk management, stochastic modeling, Monte Carlo simulation, and applied probability.

Yingda Song is an assistant professor of Antai College of Economics and Management at Shanghai Jiao Tong University. His research interests focus on financial engineering, simulation, stochastic modeling, and applied probability.

Nan Chen is an associate professor in the Department of Systems Engineering and Engineering Management at the Chinese University of Hong Kong. His research interests include quantitative methods in finance and risk management, Monte Carlo simulation, and applied probability.