Empirical Martingale Simulation for Asset Prices

Jin-Chuan Duan • Jean-Guy Simonato

Department of Finance, Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong

École des Hautes Etudes Commerciales, 3000, Chemin de la Côte-Sainte-Catherine, Montréal, Quebec, Canada H3T 2A7

This paper proposes a simple modification to the standard Monte Carlo simulation procedure for computing the prices of derivative securities. The modification imposes the martingale property on the simulated sample paths of the underlying asset price. This procedure is referred to as the empirical martingale simulation (EMS). The EMS ensures that the price estimated by simulation satisfies the rational option pricing bounds. The EMS yields a substantial error reduction for the price estimate and can be easily coupled with the standard variance reduction methods. Simulation studies are conducted for European and Asian call options using both the Black and Scholes and GARCH option pricing frameworks. The results indicate that the EMS yields substantial variance reduction particularly for in- and at-the-money or longer-maturity options. The option price estimate based on the EMS is found to exhibit a minor small-sample bias only in few occasions. An analysis of the trade-off between computing time and price accuracy reveals that the EMS dominates the conventional simulation methods.

(*Martingale*; *Option Pricing*; *Monte Carlo Simulation*; *GARCH*; *Asian Options*)

1. Introduction

Monte Carlo simulation is a widely used tool for estimating derivative security prices when there is no closed-form solution. It was first introduced by Boyle (1977) to option pricing. Monte Carlo method is especially useful when one deals with path dependent asset prices and / or option payoffs; for example, it has been used in Duan (1995) to compute the GARCH option prices (path-dependent asset prices) and in Kemna and Vorst (1990) to compute Asian option prices (path-dependent payoffs). Boyle et al. (1997) provides a comprehensive survey on the recent developments in Monte Carlo methods for option pricing.

The price of a derivative contract in an arbitrage-free economy can be expressed as a discounted expected value of its random payoffs. Monte Carlo simulation is hence a natural tool for approximating this expectation by the sample average. The commonly used Monte Carlo simulation procedure for option pricing can be briefly described as follows: first, simulate sample paths for the underlying asset price; secondly, compute its corresponding option payoff for each sample path; and finally, average the simulated payoffs and discount the average to yield the Monte Carlo price of an option. Although arbitrary degree of accuracy can in principle be achieved, Monte Carlo simulation tends to be a rather numerically intensive method if a high degree of accuracy is desired. This is, of course, due to the well-known fact that the standard error of a Monte Carlo estimate is inversely proportional to the square root of the number of simulated sample paths. A less known difficulty related to the use of Monte Carlo simulation is the occurrence of the simulated price violating rational option-pricing bounds and, hence, being a nonsensible price estimate. This bound violation could have serious implications. The implied volatility based on the Black and Scholes (1973) formula is often used as a standardized measure for examining more complex option-pricing models: for example, Hull and White (1987) and Duan (1995). When the option-pricing bound is violated, the Black-Scholes implied volatility cannot even be computed.

0025-1909 / 97 / 4409 / 1218$05.00

The idea of this paper is based on a simple observation that simulated sample paths for the underlying asset price almost always fail to possess the martingale property even though the theoretical model does. This results from the fact that simulation can only approximate the theoretical properties because of finite repetitions and the quality of the random number generator. The failure to ensure the martingale property has particularly serious consequences because the asset price dynamics are typically modeled as an exponential (semi)martingale. This multiplicative system has a rather undesirable error propagation property. It often requires a very large number of simulation repetitions to dampen simulation errors. We propose a simple correction to the standard procedure by ensuring that the simulated sample paths are together a martingale in an *empirical* sense. This correction will be referred to as empirical martingale simulation (EMS). With the EMS, one can be certain that option-pricing bounds, if they are due to Jensen's inequality, are satisfied in simulation. Apart from yielding more sensible price estimates, the use of the EMS can substantially reduce Monte Carlo errors, particularly significant for in- and at-the-money or longer-maturity options. The error reduction is irrespective of the number of sample paths and is obtained for the plain-vanilla European options, as well as for the path-dependent ones such as Asian options. The results also hold true for both the Black-Scholes and GARCH option-pricing frameworks.

The EMS can also be coupled with the standard variance reduction techniques. A number of such techniques are available for option pricing. Antithetic and control-variate simulations are perhaps the most widely known procedures. We show in this paper that the EMS can be easily incorporated into these variance reduction techniques. The EMS correction is truly simple and practically requires no additional programming efforts.

We show theoretically that the EMS yields consistent price estimates for contingent contracts, although these price estimates are likely to be biased in small samples. Our numerical analysis, however, suggests that the small-sample biases are minor and typically insignificant by the standard statistical measure. We also conduct a computational efficiency analysis comparing the EMS with the conventional Monte Carlo methods. The comparison is based on an analysis, similar to Broadie and Detemple (1996), that trades off computational time with price accuracy for a pool of option contracts. The EMS clearly dominates in terms of computational efficiency.

2. The Martingale Property in Monte Carlo Simulation

The theoretical works for contingent claim pricing mostly rely on absence of arbitrage. In Black and Scholes (1973) and Merton (1973), option-pricing formulas were derived from this principle. The martingale connection to the arbitrage-free price system was first observed by Cox and Ross (1976) and later formalized by Harrison and Kreps (1979) and Harrison and Pliska (1981). In an explicit equilibrium setting such as Lucas (1978), the martingale connection can also be established; for example, Duan (1995) used this approach to derive the GARCH option-pricing model. In this paper we explore the (semi-)martingale property in Monte Carlo simulation. For the ease of exposition, we consider a price system consisting of two securities—one risky and one risk-free. The risky security, say a common stock, does not pay dividends, and its price, denoted by $S(t)$, has the following dynamics under the risk neutral probability measure Q:

$$S(t) = S_0 \exp\left[rt - \tfrac{1}{2}\int_0^t \sigma(s)^2 ds + \int_0^t \sigma(s) dW(s) \right], \quad (1)$$

where r is the continuously compounded return on the risk-free security, $\sigma(s)$ is the instantaneous standard deviation of the asset return (satisfying some regularity conditions) and $W(s)$ is a standard Brownian motion under probability measure Q.[1] It is easy to verify that the discounted asset price is indeed a Q-martingale; that is, for any $t \geq \tau \geq 0$,

$$E^Q[e^{-rt} S(t) \,|\, \mathcal{F}_\tau] = e^{-r\tau} S(\tau), \quad (2)$$

where $E^Q(\cdot)$ denotes the expectation operator under the risk-neutral measure Q and $\mathcal{F}_\tau$ the information filtration up to time τ.

In a typical Monte Carlo simulation, this martingale property almost always fails in the simulated sample.

[1] For option pricing in the GARCH framework, a discrete-time analogous formula can be found in Duan (1995).

In other words, the discounted sample average of $S(t)$ computed from a Monte Carlo simulation will be in almost all cases different from S_0. This discrepancy between the empirical and theoretical values will, of course, depend on the sample size and the quality of the random number generator. However, even a small departure of the sample value from its theoretical one can have important consequences. For example, in the case of a European call option, the failure to *empirically* satisfy the martingale property may result in a violation of the following rational option-pricing bound:

$$C_0(t) \geq \max(S_0 - Ke^{-rt}, 0), \tag{3}$$

where K is the exercise price and $C_0(t)$ is the current price of the European call option maturing at time t.[2] This bound can be derived from a direct application of Jensen's inequality to the theoretical expression:

$$C_0(t) = e^{-rt}E^Q[\max(S(t) - K, 0) | \mathcal{F}_0]. \tag{4}$$

The sample equivalent to the right-hand side (RHS) of (4) may be less than the RHS of (3) if the discounted sample average of $S(t)$ is smaller than S_0.

This point can be better understood by considering a deep in-the-money call. Since the ''max'' function causes very few truncations at zero, the theoretical call price is well approximated by

$$C_0(t) \approx e^{-rt}[E^Q(S(t) | \mathcal{F}_0) - K] \quad \text{if } K/S_0 \text{ is small.} \tag{5}$$

Define the discounted sample average of the simulated asset prices at time t with n sample paths as

$$\hat{S}_0(t, n) = \frac{1}{n} e^{-rt} \sum_{i=1}^{n} \hat{S}_i(t), \tag{6}$$

where $\hat{S}_i(t)$ is the ith simulated asset price at time t, for $i = 1, \cdots n$. The corresponding call price by Monte Carlo simulation, denoted by $\hat{C}_0(t, n)$, must be well-approximated by

$$\hat{C}_0(t, n) = \frac{1}{n} e^{-rt} \sum_{i=1}^{n} \max[\hat{S}_i(t) - K, 0] \tag{7}$$

$$\approx \hat{S}_0(t, n) - Ke^{-rt} \quad \text{if } K/S_0 \text{ is small.} \tag{8}$$

If $\hat{S}_0(t, n)$ is smaller than S_0, it is possible that the computed call price violates the rational option-pricing bound; that is, $\hat{C}_0(t, n) < \max(S_0 - Ke^{-rt}, 0)$.

To illustrate this bound violation, we perform a simulation study using the Black and Scholes (1973) model. For this model, the asset price process is a special case of (1) with a constant volatility. The Monte Carlo option price for a given n and t can be computed according to (7). The numbers in Table 1 are the percentages of bound violation in 1,000 random repetitions. The first column of the table indicates the asset-to-strike price ratio. Three maturities—1, 3, and 9 months—are considered. For each repetition, we compute the option prices using 1,000 and 10,000 Monte Carlo sample paths, respectively. Note that the use of the Black-Scholes model implies that the bound in (3) must be a strict inequality.

The bound violation occurs for in- and out-of-the-money options. The percentage of bound violations can attain 50% for in-the-money and 90% for out-of-the-money, one-month options. The bound violation lessens when the number of sample paths increases. The occurrence of bound violations also depends on the maturity of an option. For shorter-maturity options, it is more likely to experience a bound violation even with 10,000 sample paths. These results are expected because the price of an option approaches the pricing bound when the strike price is either increased or decreased. In other words, the time-value component of an option decreases when the strike price is pushed to the extremes. When the time-value component drops in magnitude, a small Monte Carlo error can cause a bound violation. The time-value component argument also works in the dimension of option maturity and the number of sample paths. Increasing an option's maturity yields a higher time-value component, and hence causes less violations. An increase in the number of sample paths reduces Monte Carlo errors. As a result, a given size of the time-value component is more likely to be large enough to ensure no option bound violation.

3. The Empirical Martingale Simulation

In this section we present a simple correction to the standard Monte Carlo simulation procedure. This proposed correction imposes the martingale property on the

[2] The failure to *empirically* possess the martingale property can also lead to the violation of the put-call parity, among others.

Table 1 European Call Options in the Black and Scholes Framework: Percentage of Bound Violations in a Crude Monte Carlo Simulation Experiment With 1,000 Random Repetitions

	1,000 Sample Paths			10,000 Sample Paths		
S_0/K	$t = 1$ month	$t = 3$ months	$t = 9$ months	$t = 1$ month	$t = 3$ months	$t = 9$ months
1.20	51	43	18	50	28	0
1.10	34	3	0	8	0	0
1.00	0	0	0	0	0	0
0.90	0	0	0	0	0	0
0.80	92	0	0	47	0	0

Parameters: $S_0 = 100$, $r = 0.10$ (annualized) and $\sigma = 0.20$ (annualized).

collection of the simulated sample paths. Let $t_0 = 0$, the current time. This new simulation procedure, referred to as the empirical martingale simulation (EMS), generates the EMS asset prices at a sequence of future time points, $t_1, t_2, \ldots, t_m$, using the following system:

$$S_i^*(t_j, n) = S_0 \frac{Z_i(t_j, n)}{Z_0(t_j, n)}, \quad (9)$$

where

$$Z_i(t_j, n) = S_i^*(t_{j-1}, n) \frac{\hat{S}_i(t_j)}{\hat{S}_i(t_{j-1})}, \quad (10)$$

$$Z_0(t_j, n) = \frac{1}{n} e^{-rt_j} \sum_{i=1}^{n} Z_i(t_j, n). \quad (11)$$

Note that $\hat{S}_i(t)$ is the ith simulated asset price at time t prior to the EMS adjustment, and $\hat{S}_i(t_0)$ and $S_i^*(t_0, n)$ are set equal to S_0. The adjustment steps can be understood as follows. First, we take the standard simulated return from t_{j-1} to t_j, i.e., $\hat{S}_i(t_j)/\hat{S}_i(t_{j-1})$, to create a temporary asset price at time t_j, i.e., $Z_i(t_j, n)$. Second, we compute the discounted sample average, $Z_0(t_j, n)$. Finally, we compute the EMS asset price at time t_j by (9). After the EMS correction, the simulation moves on to the next time point, and repeats the whole process again.

We describe the EMS correction as a recursive scheme because the option payoff or the underlying asset price dynamics may be path-dependent so that the simulation must be conducted recursively until the option's maturity. The desirable frequency of the EMS correction entirely depends on the nature of the problem; for example, an Asian option based on the path average of the daily prices will call for the EMS correction to be performed on a daily recursive basis. Note that the EMS conducts a simulation of n sample points at one time, which is required for making the EMS correction. If the use of a very large n causes a memory allocation problem, the sample can be broken into batches with each undergoing the EMS correction separately and being combined in the end.

The EMS bears some resemblance to the moment-matching simulation (MMS) of Barraquand (1995).[3] The EMS is, however, entirely different from the MMS in the following sense. For the EMS, the correct first moment in simulation is ensured by using a multiplicative adjustor, instead of using an additive one as in the moment-matching adjustment if it is applied to the asset price directly. This difference is extremely important because asset prices are typically modeled as exponential (semi-)martingales. The multiplicative adjustment ensures no domain violation, whereas the additive adjustment cannot.[4] If the moment-matching adjustment is applied to the exponent of an exponential martingale (logarithmic asset price), in which case there will be no domain violation, the simulated asset prices will no longer be an empirical martingale, however. The failure to maintain the martingale property may have important consequences, especially when one needs to simulate asset prices at multiple time points. In this case, the

[3] We use the terminology of Boyle et al. (1997) to refer to the quadratic resampling method of Barraquand (1995) as the moment-matching simulation.

[4] Footnote 7 in Boyle et al. (1997) provides a discussion on domain violation of Barraquand's (1995) moment-matching simulation.

failure to ensure the martingale property over the subperiods may produce an error propagation that leads to a substantial departure from the theoretical model over the option's maturity. Since the EMS ensures the martingale property for all time points of the simulation, it thus precludes this possibility of error propagation. In short, the EMS correction is a more natural adjustment, which is made to the entire simulated sample paths as required by the option-pricing theory.

The discounted EMS asset price estimate at time 0 is therefore

$$S_0^*(t, n) = \frac{1}{n} e^{-rt} \sum_{i=1}^{n} S_i^*(t, n) \tag{12}$$

$$= S_0. \tag{13}$$

This is true for any n and $t_1, t_2, \ldots, t_m$. Since the EMS asset price can be regarded as an empirical martingale with equal probabilities assigned to all simulated sample paths, any option-pricing bound, as long as it results from Jensen's inequality, is always satisfied. Formally, we define the EMS option price estimate as

$$C_0^*(t, n) = \frac{1}{n} e^{-rt} \sum_{i=1}^{n} \max[S_i^*(t, n) - K, 0]. \tag{14}$$

It follows that $C_0^*(t, n) \geq \max[S_0^*(t, n) - Ke^{-rt}, 0] = \max(S_0 - Ke^{-rt}, 0)$.

The EMS for derivatives pricing also preserves consistency, i.e., the convergence to the theoretical value, under fairly general conditions. This is a relevant because the EMS sample paths are by construction not independent. The following proposition provides a general characterization for consistency.

PROPOSITION 1. *Let $e^{-rt_j}S(t_j)$ be a positive Q-martingale over the time index set $\{t_j : j = 0, \ldots, m\}$. If the payoff function, $f[S(t_1), \ldots, S(t_m)]$ from R_+^m to R, is Lipschitz continuous, i.e., there exists $c < \infty$ such that $|f(x) - f(y)| \leq c\|x - y\|$ for any $x \in R_+^m$ and $y \in R_+^m$ where $\|\cdot\|$ stands for the Euclidean norm, then as $n \to \infty$,*

$$\frac{1}{n} \sum_{i=1}^{n} f[S_i^*(t_1, n), \ldots, S_i^*(t_m, n)] \to E^Q\{f[S(t_1), \ldots, S(t_m)] \mid \mathcal{F}_0\} \text{ almost surely.} \tag{15}$$

PROOF. See Appendix.

Proposition 1 states that for a well-behaved payoff function, the EMS estimate converges to its theoretical value. For a European option, it is clear that the payoff function is Lipschitz continuous. The EMS price estimate is thus consistent irrespective of the pricing framework being the Black and Scholes, GARCH, or anything else. Most exotic options, such as Asian options and lookback options, also have Lipschitz continuous payoff functions; Proposition 1 thus applies to many derivative contracts.

Proposition 1 is limited in the sense that digital (binary) options are excluded because they have discontinuous payoff functions. This type of contract can nevertheless be dealt with by strengthening the result in Proposition 1. To proceed with the discussion, we need the following definition.

DEFINITION 1. A function $f(x)$, mapping from R_+^m to R, is said to satisfy the generic Lipschitz condition if there exists $b < \infty$ such that $|f(x)| < b(1 + \|x\|)$ for any $x \in R_+^m$ where $\|\cdot\|$ stands for the Euclidean norm, and there exists a finite partition of its domain such that every element of the partition is a connected set and the function is Lipschitz continuous over any element of the partition; i.e., there exists $c < \infty$ such that $|f(x) - f(y)| \leq c\|x - y\|$ for any $x \in A$ and $y \in A$ where A is any element of the partition.

REMARK. 1. The constant c need not depend on the element of the partition because the partition is finite.

2. The condition, $|f(x)| < b(1 + \|x\|)$, is used to bound the function's value at the points of discontinuity.

PROPOSITION 2. *Let $e^{-rt_j}S(t_j)$ be a positive Q-martingale over the time index set $\{t_j : j = 0, \ldots, m\}$. If the payoff function, $f[S(t_1), \ldots, S(t_m)]$, satisfies the generic Lipschitz condition and the multivariate distribution for $[S(t_1), \ldots, S(t_m)]$ has a bounded density function, then as $n \to \infty$,*

$$\frac{1}{n} \sum_{i=1}^{n} f[S_i^*(t_1, n), \ldots, S_i^*(t_m, n)] \to E^Q\{f[S(t_1), \ldots, S(t_m)] \mid \mathcal{F}_0\} \text{ almost surely.} \tag{16}$$

PROOF. See Appendix.

The result in Proposition 2 is a generalization in terms of the payoff function, and this generalization is obtained by imposing a restriction on distributions. The payoff function of a digital option takes on two values—for example, 1 if $S_t \geq K$ and 0 if $S_t < K$. Although the payoff function is discontinuous, it satisfies the generic Lipschitz condition if one considers the partition $\{[0, K), [K, \infty)\}$. The EMS price estimates for digital options must therefore be consistent under the usual distributional assumption.

Another attractive feature of the EMS is its ability to reduce the Monte Carlo simulation error. The source of error reduction comes from the fact that the EMS ensures the martingale property in the simulated sample. Because the discounted sample average asset price, $S_0^*(t, n)$, always equals S_0, $\mathrm{Var}[S_0^*(t, n)] = 0$. This is, however, not true for $\hat{S}_0(t, n)$, i.e., $\mathrm{Var}[\hat{S}_0(t, n)] > 0$. The implication of this property for option prices is best seen for in-the-money options. The call prices computed using the Monte Carlo simulation with and without the EMS correction can be written approximately as

$$\hat{C}_0(t, n) \approx \hat{S}_0(t, n) - Ke^{-rt} \quad \text{if } K/S_0 \text{ is small,} \tag{17}$$

$$C_0^*(t, n) \approx S_0^*(t, n) - Ke^{-rt} \quad \text{if } K/S_0 \text{ is small.} \tag{18}$$

This in turn implies that

$$\mathrm{Var}[\hat{C}_0(t, n)] \approx \mathrm{Var}[\hat{S}_0(t, n)] > 0 \quad \text{if } K/S_0 \text{ is small,} \tag{19}$$

$$\mathrm{Var}[C_0^*(t, n)] \approx \mathrm{Var}[S_0^*(t, n)] = 0 \quad \text{if } K/S_0 \text{ is small.} \tag{20}$$

The EMS option price estimate therefore has a negligible variance if the option is deep in-the-money. The numerical significance of the EMS error reduction is examined in the next section.

4. The EMS Applications

The application of the EMS to two option-pricing frameworks—Black and Scholes (1973) and GARCH (Duan, 1995)—is studied here. For each pricing framework, we consider two types of contracts—European calls and Asian calls. We choose two pricing frameworks and two different types of option contracts to demonstrate path-dependency in terms of option payoffs and/or asset price dynamics. Asian options give an interesting case of path-dependent option payoffs that require Monte Carlo simulation, whereas the GARCH framework provides a situation in which the asset price dynamics is path-dependent.

For each model and contract, we generate 500 Monte Carlo option price estimates. These prices are computed with 10,000 sample paths.[5] We consider three maturities—1, 3, and 9 months—and three asset-to-strike price ratios—1.1, 1, and 0.9. In our numerical analyses, one year is assumed to have 365 days for the annualizing purpose. We compare the EMS with the standard Monte Carlo simulation and with the MMS of Barraquand (1995).[6] We also incorporate variance reduction techniques such as antithetic and control-variate simulation whenever appropriate.

4.1. The Black and Scholes Option Pricing Framework

For European call options in the Black and Scholes (1973) framework, the underlying asset price is simulated using (1) and setting $\sigma(t)$ to a constant. This simulation can be done in one step to the option's maturity because the problem is not path dependent.[7] For the EMS, we simply take the simulated random variates and apply once the EMS correction. The results are presented in Table 2.

For European calls, the mean value of either the crude Monte Carlo or the EMS prices in 500 repetitions is very close to its theoretical price. As expected, for all cases presented in this table, the EMS option prices have smaller standard deviations when compared to the crude Monte Carlo simulation. When compared to the MMS of Barraquand (1995), the EMS also yields smaller

[5] We use the built-in normal random number generator of the matrix programming language GAUSS version 3.2.12. In the earlier versions of this paper, we have also used 1,000 sample paths to compute option prices. The error reduction ratios reported later in Tables 2–5 are insensitive to the change in sample size.

[6] In the tables presented in this section, we only consider the first-moment matching adjustment for the exponent of the stock price. Although in the Black-Scholes framework the second-moment adjustment suggested by Barraquand (1995) is possible, one can also strengthen the EMS adjustment by a second-moment correction. The second-moment adjustment is not always possible, however—for example, in the GARCH option pricing framework.

[7] For European options, a closed-form solution exists. We only use these options as a means to compare the EMS with the standard simulation procedure.

Table 2 European Call Options in the Black and Scholes Framework

	10,000 Sample Paths								
	Maturity = 1 Month			Maturity = 3 Months			Maturity = 9 Months		
S_0/K	1.10	1.00	0.90	1.10	1.00	0.90	1.10	1.00	0.90
Theoretical	9.9117	2.7104	0.1116	11.8209	5.2498	1.2147	16.9270	10.7748	5.4842
Crude									
mean	9.9133	2.7113	0.1117	11.8238	5.2517	1.2148	16.9319	10.7794	5.4838
std.	(0.0560)	(0.0366)	(0.0072)	(0.0926)	(0.0682)	(0.0345)	(0.1575)	(0.1316)	(0.1009)
EMS									
mean	9.9117	2.7098	0.1116	11.8212	5.2490	1.2148	16.9276	10.7742	5.4837
std.	(0.0052)	(0.0170)	(0.0066)	(0.0165)	(0.0287)	(0.0252)	(0.0324)	(0.0459)	(0.0528)
MMS									
mean	9.9117	2.7098	0.1116	11.8212	5.2490	1.2148	16.9275	10.7740	5.4837
std.	(0.0068)	(0.0183)	(0.0066)	(0.0220)	(0.0328)	(0.0266)	(0.0492)	(0.0599)	(0.0621)
Ratio of std.:									
Crude/EMS	10.8048	2.1505	1.0926	5.6176	2.3739	1.3666	4.8568	2.8643	1.9106
MMS/EMS	1.3077	1.0765	1.0000	1.3333	1.1429	1.0556	1.5185	1.3050	1.1761

Crude: Crude Monte Carlo simulation. EMS: Empirical martingale simulation. MMS: Moment matching simulation. EMS-AT: Antihetic empirical martingale simulation. Parameters: $S_0 = 100$, $r = 0.10$ (annualized) and $\sigma = 0.20$ (annualized). All standard errors are computed by repeating (randomly) the calculation 500 times.

standard deviations. The last two rows report the ratios of the crude simulation to the EMS and the MMS to the EMS. Changes in this ratio are observed when the maturity and exercise price are varied. For in-the-money, short-maturity options, the ratio reveals a phenomenal error reduction when compared to the crude simulation. The standard deviation of the EMS price is smaller by a factor of approximately 10. The decrease in the standard deviation is very substantial for at-the-money, short-maturity options with a factor of approximately 2. When the maturity of the option is increased, the error reduction for in-the-money options becomes smaller but is still substantial. For the EMS, this can be explained by a larger probability for the asset price to finish out-of-the-money when the maturity is increased. In short, whenever there is a larger number of truncations caused by the ''max'' function, the EMS has a smaller efficiency gain. For out-of-the-money options, the ratio of standard deviations increases as the maturity increases. This can be explained by the relative increase in the probability to finish in-the-money and, hence, a larger EMS efficiency gain. The increase in maturity therefore has an effect of leveling the error reduction across moneyness positions. The efficiency gain of the EMS is nevertheless significant across the board. The comparison between the EMS and the MMS in Table 2 shows a similar pattern, which reveals a measurable dominance of the EMS over the MMS, particularly for in-the-money or longer-maturity options.

For Asian call options, the entire sample path of the asset price becomes important because the option payoff is based on the path average of the asset prices. We assume that the path average is calculated on daily closing prices. The asset prices must then be simulated daily according to

$$S(t+1) = S(t) \exp\left[r - \frac{\sigma^2}{2} + \sigma \varepsilon_{t+1} \right], \quad (21)$$

where r and σ are the daily risk-free rate and standard deviation, respectively. The random variable ε_{t+1} is an element of an i.i.d. standard normal random sequence. For the ith sample path, the payoff of the Asian call option with maturity T equals $\max(\hat{A}_{i,T} - K, 0)$ where $\hat{A}_{i,T} = (1/T) \sum_{t=1}^{T} \hat{S}_i(t)$. The ith sample path for the EMS is obtained using the correction to $\hat{S}_i(t)$ described in (9) period-by-period to the maturity of the option. The path average and option payoff are then computed accordingly.

For the Asian option example, we also use the antithetic variance reduction technique. The antithetic simulation effectively increases the sample size by a factor of 2. This factor should be taken into consideration if one intends to compare the simulation results with and without the antithetic sample. We are only interested in comparing the antithetic Monte Carlo simulation with and without the EMS correction. The antithetic simulation with the EMS can be easily performed by applying the EMS correction to the antithetic sample. The results for Asian options are reported in Table 3.

For Asian options, the patterns are similar to those for European options. The decrease in Monte Carlo errors due to the use of the EMS is much larger for in-the-money, short-maturity options with a ratio in excess of 60. For Asian options, the use of the antithetic variable technique produces large decreases in standard deviations. With the EMS correction being added onto the antithetic variable simulation, the additional improvement is clear as indicated by the ratios of standard deviations reported in the last rows of the table. As in the case of European options, the EMS dominates the MMS, yielding a larger variance reduction if the option is in-the-money or with a longer maturity.

4.2. The GARCH Option Pricing Framework

In the GARCH option-pricing framework of Duan (1995), the asset price has a form analogous to equation (1), where the integrals are replaced by summations and the Brownian motion by the standard normal random variable. In this section, we use the GARCH(1, 1)-(in mean) model to describe the daily asset return dynamics. According to Duan (1995), the dynamics under the locally risk-neutralized probability measure Q is

Table 3 **Asian Call Options in the Black and Scholes Framework**

	10,000 Sample Paths								
	Maturity = 1 Month			Maturity = 3 Months			Maturity = 9 Months		
S_0/K	1.10	1.00	0.90	1.10	1.00	0.90	1.10	1.00	0.90
Crude									
mean	9.4419	1.5687	0.0016	10.1595	2.9352	0.1448	12.3531	5.8069	1.5078
std.	(0.0312)	(0.0211)	(0.0006)	(0.0531)	(0.0391)	(0.0083)	(0.0857)	(0.0698)	(0.0374)
EMS									
mean	9.4403	1.5683	0.0016	10.1585	2.9348	0.1447	12.3523	5.8048	1.5073
std.	(0.0005)	(0.0102)	(0.0006)	(0.0044)	(0.0170)	(0.0075)	(0.0132)	(0.0264)	(0.0274)
MMS									
mean	9.4403	1.5683	0.0016	10.1585	2.9348	0.1447	12.3520	5.8049	1.5072
std.	(0.0012)	(0.0107)	(0.0006)	(0.0065)	(0.0185)	(0.0077)	(0.0196)	(0.0315)	(0.0295)
AT									
mean	9.4403	1.5683	0.0016	10.1584	2.9350	0.1445	12.3519	5.8055	1.5075
std.	(0.0012)	(0.0107)	(0.0004)	(0.0059)	(0.0182)	(0.0061)	(0.0184)	(0.0310)	(0.0261)
EMS-AT									
mean	9.4403	1.5683	0.0016	10.1585	2.9349	0.1445	12.3521	5.8053	1.5074
std.	(0.0004)	(0.0102)	(0.0004)	(0.0034)	(0.0166)	(0.0059)	(0.0109)	(0.0258)	(0.0238)
Ratio of std.:									
Crude/EMS	62.1378	2.0724	1.0075	12.0842	2.3068	1.0983	6.5221	2.6449	1.3618
MMS/EMS	2.4000	1.0490	1.0000	1.4773	1.0882	1.0267	1.4848	1.1932	1.0766
AT/EMS-AT	3.3471	1.0490	1.0021	1.7410	1.0952	1.0280	1.6958	1.2046	1.0961

Crude: Crude Monte Carlo simulation. EMS: Empirical martingale simulation. MMS: Moment matching simulation. EMS-AT: Antihetic empirical martingale simulation. Parameters: $S_0 = 100$, $r = 0.10$ (annualized) and $\sigma = 0.20$ (annualized). All standard errors are computed by repeating (randomly) the calculation 500 times.

Table 4 **European Call Options in the GARCH Framework**

	10,000 Sample Paths								
	Maturity = 1 Month			Maturity = 3 Months			Maturity = 9 Months		
S_0/K	1.10	1.00	0.90	1.10	1.00	0.90	1.10	1.00	0.90
Crude									
mean	9.9236	2.5374	0.1164	11.7566	5.0119	1.0615	16.7574	10.4646	5.1194
std.	(0.0501)	(0.0349)	(0.0089)	(0.0830)	(0.0671)	(0.0317)	(0.1380)	(0.1219)	(0.0936)
EMS									
mean	9.9224	2.5367	0.1164	11.7573	5.0104	1.0610	16.7580	10.4650	5.1190
std.	(0.0067)	(0.0176)	(0.0084)	(0.0168)	(0.0281)	(0.0248)	(0.0311)	(0.0452)	(0.0527)
MMS									
mean	9.9229	2.5368	0.1163	11.7568	5.0095	1.0605	16.7582	10.4627	5.1181
std.	(0.0140)	(0.0202)	(0.0086)	(0.0302)	(0.0352)	(0.0272)	(0.0625)	(0.0654)	(0.0664)
CV									
mean	9.9226	2.5363	0.1163	11.7578	5.0106	1.0615	16.7573	10.4625	5.1186
std.	(0.0134)	(0.0109)	(0.0058)	(0.0255)	(0.0214)	(0.0152)	(0.0462)	(0.0407)	(0.0364)
EMS-CV									
mean	9.9222	2.5361	0.1163	11.7578	5.0111	1.0618	16.7571	10.4644	5.1189
std.	(0.0044)	(0.0070)	(0.0056)	(0.0089)	(0.0125)	(0.0123)	(0.0146)	(0.0195)	(0.0247)
Ratio of std.:									
Crude/EMS	7.4665	1.9799	1.0583	4.9382	2.3855	1.2767	4.4387	2.6998	1.7767
MMS/EMS	2.0896	1.1477	1.0238	1.7976	1.2527	1.0968	2.0096	1.4469	1.2600
CV/EMS-CV	3.0478	1.5700	1.0382	2.8603	1.7103	1.2414	3.1613	2.0887	1.4736

Crude: Crude Monte Carlo simulation. EMS: Empirical martingale simulation. MMS: Moment matching simulation. EMS-CV: Control variate empirical martingale simulation. Parameters: $S_0 = 100$, $r = 0.10$ (annualized), $\beta_0 = 0.00001$, $\beta_1 = 0.70$, $\beta_2 = 0.20$ and $\lambda = 0.01$. All standard errors are computed by repeating (randomly) the calculation 500 times.

$$\ln \frac{S(t+1)}{S(t)} = r - \frac{1}{2}\sigma(t+1)^2 + \sigma(t+1)\varepsilon_{t+1}, \quad (22)$$

$$\sigma(t+1)^2 = \beta_0 + \beta_1\sigma(t)^2 + \beta_2\sigma(t)^2(\varepsilon_t - \lambda)^2, \quad (23)$$

$$\varepsilon_{t+1} | F_t \overset{Q}{\sim} N(0, 1), \quad (24)$$

where β_0, β_1, β_2 are the GARCH(1, 1) parameters, and λ is the unit risk premium (per unit of conditional standard deviation) parameter that defines the conditional mean equation of the GARCH(1, 1)-(in mean) dynamics under the physical probability measure. By Duan (1995), the stationary variance of the daily asset return under the locally risk-neutralized probability measure Q is $\beta_0[1 - \beta_1 - \beta_2(1 + \lambda^2)]^{-1}$, which is higher than $\beta_0(1 - \beta_1 - \beta_2)^{-1}$, the stationary variance under the physical probability measure.

In our numerical study, we set the initial conditional variance at the stationary level under the physical probability measure, i.e., $\sigma(1)^2 = \beta_0(1 - \beta_1 - \beta_2)^{-1}$. The crude Monte Carlo simulation is conducted according to the GARCH dynamics under Q described above. The EMS is conducted by applying the EMS correction described in (9) to the simulated sample paths. For European options, we use their Black–Scholes counterparts in the control-variate Monte Carlo simulation. Following the practice in Duan (1995), the Black-Scholes price is computed using the stationary variance $\beta_0(1 - \beta_1 - \beta_2)^{-1}$ in the Black-Scholes closed-form formula. The results for European call options are presented in Table 4. The error reductions due to the EMS are roughly the same as those under the Black-Scholes framework reported in Table 2. The use of control-variate simulation reduces simulation errors substantially. If one uses the EMS in conjunction with the control-variate simulation, it will further improve simulation efficiency in a significant way. The efficiency gain is most pronounced for in-the-money options. For out-of-the-money options, the efficiency gain increases as the maturity increases. The comparison of the EMS with the MMS again shows that the EMS dominates in terms of variance reduction.

For Asian options, we use the contract specification described in the previous subsection. To conduct the control-variate simulation, we have available the closed-form formula for geometric average options developed by Turnbull and Wakeman (1991).[8] The results in Table 5 again clearly show the improvement due to the EMS. The patterns are the same as the ones for European options reported in Table 4.

5. A Performance Assessment of the EMS

5.1. Small-Sample Biases

We have proved earlier that the EMS option price estimator is consistent. For all derivative securities with linear payoffs, the EMS price estimator will have no pricing error for any finite sample, since it is based on an adjustment ensuring that the simulated sample average price always equals its theoretical value. However, for a derivative security with nonlinear payoffs, such as options, the EMS price estimator is expected to be biased in a finite sample. We examine this possibility in this subsection.

To get some ideas with respect to the EMS's performance when small simulation samples are used, we conduct the following analysis. We consider European options in the Black-Scholes framework. This case facilitates the analysis because the option's theoretical value can be analytically computed. We vary the size of the simulation sample from 100 to 1,000. As noted earlier, the EMS adjustment destroys the independence of the simulated sample. We must therefore repeat the calculation many times in order to draw a statistical inference. For this purpose, 1,000 random repetitions are used. Consider a case for which the number of sample paths equals 100. We compute the option price using the EMS with 100 sample paths, which yields one EMS option price. Randomly repeating 1,000 times this calculation gives rise to a sample of 1,000 independent EMS option prices. We can infer from this random sample as to whether the EMS option price is biased. Specifically, if the EMS option price is unbiased, the difference between the average of the 1,000 EMS option prices and the theoretical option price must asymptotically distribute as a normal random variable whose mean equals zero and variance is approximated by the sample variance divided by 1,000.

The results for different maturities and exercise prices are reported in Table 6. It is clear that statistically significant biases only occasionally occur. In short, we can conclude that the EMS does not produce statistically significant biases when the simulated sample is small.

5.2. Computational Efficiency

In this subsection, we perform a cost-benefit analysis for different Monte Carlo methods in a manner similar to the procedure adopted by Broadie and Detemple (1996). The comparison analysis begins by choosing a large test pool of options using a random selection of parameters based on predetermined distributions. For each Monte Carlo method, we record the total computing time and an aggregate measure of pricing errors for the test pool of options.

The comparison analysis is only conducted for European options in the Black and Scholes framework because the closed-form solution makes it easy in computing pricing errors. As in Broadie and Detemple (1996), the following distributions are used for the parameter values: σ is distributed uniformly between 0.1 and 0.6; T is also distributed uniformly between 0.1 and 1.0 year with a probability of 0.75 and uniformly between 1.0 and 5.0 years with a probability of 0.25; $S_0 = 100$ and K is uniform between 70 and 130; r is, with probability 0.8, uniform between 0.0 and 0.10, and with probability 0.2, equal to 0.0. Each parameter value is drawn independently of the others. The computing time is measured by the total number of seconds taken to complete the calculations for the test pool of options using the matrix programming language GAUSS version 3.1.5 and a standard desk-top computer equipped with a 133-MHz Pentium processor. The aggregate pricing error measure is the standard root mean squared error

[8] The same formula can also be found in Ritchken et al. (1993). Our method is similar to that of Kemna and Vorst (1990) except that we use the pricing formula for the geometric average of discretely-sampled prices, instead of continuously-sampled prices, to avoid biases.

Table 5 Asian Call Options in the GARCH Framework

	10,000 Sample Paths								
	Maturity = 1 Month			Maturity = 3 Months			Maturity = 9 Months		
S_0/K	1.10	1.00	0.90	1.10	1.00	0.90	1.10	1.00	0.90
Crude									
mean	9.4470	1.4651	0.0070	10.1645	2.7853	0.1366	12.3036	5.6098	1.3434
std.	(0.0297)	(0.0201)	(0.0020)	(0.0503)	(0.0372)	(0.0094)	(0.0823)	(0.0667)	(0.0351)
EMS									
mean	9.4451	1.4647	0.0070	10.1632	2.7850	0.1365	12.3027	5.6088	1.3430
std.	(0.0016)	(0.0102)	(0.0019)	(0.0056)	(0.0166)	(0.0087)	(0.0134)	(0.0261)	(0.0263)
MMS									
mean	9.4455	1.4646	0.0070	10.1635	2.7848	0.1364	12.3025	5.6076	1.3425
std.	(0.0072)	(0.0112)	(0.0019)	(0.0149)	(0.0200)	(0.0090)	(0.0300)	(0.0345)	(0.0296)
CV									
mean	9.4455	1.4642	0.0070	10.1635	2.7840	0.1366	12.3030	5.6080	1.3441
std.	(0.0072)	(0.0053)	(0.0018)	(0.0142)	(0.0116)	(0.0059)	(0.0255)	(0.0215)	(0.0155)
EMS-CV									
mean	9.4451	1.4642	0.0070	10.1632	2.7840	0.1366	12.3028	5.6089	1.3442
std.	(0.0014)	(0.0036)	(0.0017)	(0.0038)	(0.0070)	(0.0056)	(0.0078)	(0.0117)	(0.0121)
Ratio of std.:									
Crude/EMS	18.1469	1.9648	1.0110	9.0171	2.2340	1.0748	6.1458	2.5542	1.3362
MMS/EMS	4.5000	1.0980	1.0000	2.6607	1.2048	1.0345	2.2388	1.3218	1.1255
CV/EMS-CV	4.9920	1.4793	1.0098	3.7166	1.6496	1.0587	3.2581	1.8368	1.2780

Crude: Crude Monte Carlo simulation. EMS: Empirical martingale simulation. MMS: Moment matching simulation. EMS-CV: Control variate empirical martingale simulation. Parameters: $S_0 = 100$, $r = 0.10$ (annualized), $\beta_0 = 0.00001$, $\beta_1 = 0.70$, $\beta_2 = 0.20$ and $\lambda = 0.01$. All standard errors are computed by repeating (randomly) the calculation 500 times.

$$\mathrm{RMS}(m, n) = \sqrt{\frac{1}{m}\sum_{i=1}^{m} e_i^2(n)}\,, \tag{25}$$

where $e_i(n) = |\hat{C}_i(n) - C_i| / C_i$ with C_i the ith analytical Black and Scholes price and $\hat{C}_i(n)$ is the ith estimated option price using one of the Monte Carlo methods with n sample paths. The variable m stands for the size of the test pool. In our comparison study, we use $m = 5{,}000$ and eliminate the cases where $C_i < 0.50$ to avoid large relative errors caused by a small divider. We control the test pool to be the same for different Monte Carlo methods using common random numbers. In the Monte Carlo calculation of option prices, different methods also share a common set of random numbers. For each value of n, we produce a pair—the computing time and the relative pricing error. Figure 1 is generated by varying the value of n from 100 to 1,000 with an increment of 100.

The results are summarized in Figure 1 in which six curves are presented. The horizontal axis marks *RMS* whereas the vertical axis gives the total computing time in seconds for the test pool of options. Each curve depicts the performance of a specific Monte Carlo method using ten points generated by different values of n. If a curve is closer to the origin, it implies a better computational efficiency, because for a given level of *RMS*, it takes less time to complete the calculations. It is clear from Figure 1 that the straightforward EMS dominates other Monte Carlo methods. It is interesting to observe the performance improvement of the antithetic variable approach over the crude Monte Carlo simulation. Adding the antithetic simulation onto the MMS or EMS, however, worsens the computational efficiency, which means the marginal gain in price accuracy due to the antithetic variable is more than offset by the additional computing time required.

Table 6 Small-Sample Properties of EMS (European Call Options In the Black and Scholes Framework)

	Maturity = 1 Month			Maturity = 3 Months			Maturity = 9 Months		
S_0/K	1.10	1.00	0.90	1.10	1.00	0.90	1.10	1.00	0.90
Theoretical	9.9117	2.7104	0.1116	11.8209	5.2498	1.2147	16.9270	10.7748	5.4842
100 paths									
mean	9.9094	2.7028	0.1099	11.8131	5.2379	1.2016	16.9099	10.7598	5.4436
z-stat.	(−1.4469)	(−1.4253)	(−0.8103)	(−1.5715)	(−1.3386)	(−1.6430)	(−1.7385)	(−1.0574)	(−2.4140)
200 paths									
mean	9.9112	2.6994	0.1101	11.8201	5.2314	1.2094	16.9272	10.7468	5.4708
z-stat.	(−0.4467)	(−2.8794)	(−1.0090)	(−0.2116)	(−2.8456)	(−0.9437)	(0.0322)	(−2.7415)	(−1.1584)
300 paths									
mean	9.9116	2.7128	0.1095	11.8202	5.2547	1.2075	16.9258	10.7843	5.4686
z-stat.	(−0.0750)	(0.7800)	(−1.8290)	(−0.2464)	(0.9407)	(−1.5999)	(−0.2068)	(1.1491)	(−1.6447)
400 paths									
mean	9.9117	2.7107	0.1102	11.8218	5.2505	1.2074	16.9286	10.7765	5.4688
z-stat.	(0.0134)	(0.1167)	(−1.3128)	(0.3589)	(0.1599)	(−1.8234)	(0.3391)	(0.2276)	(−1.8531)
500 paths									
mean	9.9112	2.7047	0.1111	11.8207	5.2411	1.2099	16.9278	10.7627	5.4750
z-stat.	(−0.6204)	(−2.4111)	(−0.4845)	(−0.0522)	(−2.1962)	(−1.3312)	(0.1844)	(−1.9142)	(−1.2450)
1,000 paths									
mean	9.9122	2.7081	0.1116	11.8209	5.2461	1.2150	16.9260	10.7694	5.4829
z-stat.	(1.0442)	(−1.3848)	(−0.0368)	(0.0447)	(−1.2919)	(0.1237)	(−0.3005)	(−1.1864)	(−0.2605)

EMS: Empirical martingale simulation. Parameters: $S_0 = 100$, $r = 0.10$ (annualized) and $\sigma = 0.20$ (annualized). All z-statistics (comparing the mean and its corresponding theoretical value) are computed by repeating (randomly) the calculation 1,000 times.

Figure 1 Computational Efficiency

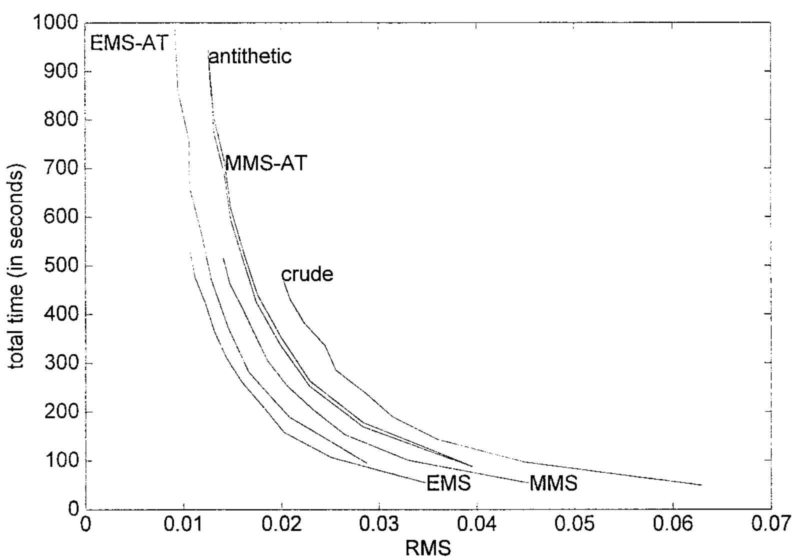

6. Conclusion

Asset prices are typically modeled as exponential (semi-) martingales. The prices of its derivative contracts are often complex functionals that sometimes require the use of Monte Carlo simulation to compute their values. Although arbitrary degree of accuracy can be obtained by simulation, it often requires too many Monte Carlo repetitions. The Monte Carlo error can also cause the price estimate to violate rational option-pricing bounds, which makes these price estimates nonsensible. We propose a simple modification to the standard Monte Carlo simulation procedure. This modification is referred to as the empirical martingale simulation (EMS), because it imposes upon the simulated sample a martingale property. The EMS ensures that the price estimated by simulation satisfies rational option-pricing bounds. Since the EMS reproduces a key feature of the theoretical model in simulation, it yields a substantial reduction in Monte Carlo errors. The EMS is applied to European and Asian call options under the Black–Scholes and GARCH frameworks to study the effect of path-dependency in payoffs and / or asset price dynamics. Substantial error reduction is obtained for all contracts under the two modeling approaches. The error reduction is particularly pronounced for in- and at-the-money or longer-maturity options. Our computational efficiency analysis, trading off computing time with price accuracy, also reveals that the EMS dominates other Monte Carlo methods.

The benefit of using the EMS goes beyond Monte Carlo error reduction. The Black-Scholes formula is often used to compute implied volatility which is a means of comparing more complex option-pricing models for options with different strike prices and maturities. For example, one may want to study the validity of a stochastic volatility option-pricing model by comparing the implied volatilities computed from the market prices of exchange-traded options and their corresponding model prices. In order to compute the implied volatility, the option price estimate must satisfy rational option-pricing bounds. The EMS is a simple way of ensuring that the rational option-pricing bounds (implied by Jensen's inequality) will be satisfied by the Monte Carlo price estimate.

A minor drawback of the EMS is related to its inability to provide a standard error estimate of its Monte Carlo price from using only one simulated sample. The efficiency gain by the EMS is derived from destroying the independence of the simulated sample paths. Because of its lack of independence, one cannot directly infer the standard error for the price estimate from one simulated sample. It also prevents employment of sequential statistical techniques for adaptively stopping the simulation when a desired accuracy has been achieved.[9]

[9] Duan is at Hong Kong University of Science and Technology, and Simonato is at HEC (Montreal) and CIRANO. This research is supported by SSHRC of Canada, grant # 410-95-0734. The earlier version was presented at HEC (Montreal), Laval University, National Central University, University of Quebec in Montreal, Hong Kong University of Science and Technology, 1996 Derivative Security Symposium at Queen's University, 1996 Japanese Association of Financial Econometrics and Engineering Conference, 1996 EIASM Finance Workshop, 1996 Mathematical Finance Conference at University of Aarhus, and 1996 Northern Finance Association Conference. The comments received from the participants are gratefully acknowledged. We also want to thank Phelim Boyle, Andrew Carverhill, Freddy Delbaen, Geneviève Gauthier, Allanus Tsoi, and three referees of this journal for their valuable inputs.

Appendix

Proof of Proposition 1

Consider the ith simulated sample path, $[\hat{S}_i(t_1), \ldots, \hat{S}_i(t_m)]$. The pre- and post-EMS adjustment values—$\hat{S}_i(t_j)$ and $S_i^*(t_j, n)$—differ by a multiplicative factor $a(n, j) = \Pi_{l=1}^{j}[Z_0(t_l, n) / S_0]$, which can be derived by repeatedly applying (9), (10), and (11). We first show that $a(n, j)$ converges to 1 almost surely as $n \to \infty$. By equations (10) and (11),

$$Z_0(t_1, n) = \frac{1}{n} e^{-rt_1} \sum_{i=1}^{n} \hat{S}_i(t_1) = \hat{S}_0(t_1, n).$$

By the martingale assumption, $e^{-rt_1}E^Q\{\hat{S}_i(t_1) | \mathcal{F}_0\} = S_0$. Because of random sampling, $\hat{S}_i(t_1)$ for $i = 1, 2, \cdots$ form an *i.i.d.* sequence. Ergodic theorem implies that $\hat{S}_0(t_1, n)$ converges to S_0 almost surely as $n \to \infty$. It also implies by equations (9) and (10) that $S_i^*(t_1, n)$ converges almost surely to $\hat{S}_i(t_1)$ for all i's. For t_2,

$$Z_0(t_2, n) = \frac{1}{n} e^{-rt_2} \sum_{i=1}^{n} \frac{S_i^*(t_1, n)}{\hat{S}_i(t_1)} \hat{S}_i(t_2) \to S_0$$

almost surely as $n \to \infty$. This is due to two reasons. First, $S_i^*(t_1, n) / \hat{S}_i(t_1) \to 1$ almost surely as $n \to \infty$ for all i's, a fact established earlier. Second,

$$\frac{1}{n} e^{-rt_2} \sum_{i=1}^{n} \hat{S}_i(t_2) = \hat{S}_0(t_2, n) \to S_0$$

almost surely as $n \to \infty$, again a result of the martingale assumption and ergodic theorem. Furthermore, by equations (9) and (10),

$$S_i^*(t_2, n) = \frac{S_0}{Z_0(t_2, n)} \frac{S_i^*(t_1, n)}{\hat{S}_i(t_1)} \hat{S}_i(t_2),$$

which converges almost surely to $\hat{S}_i(t_2)$ for all i's because both ratios in the expression converge to 1 almost surely. Repeating the same argument to yield $(Z_0(t_l, n)/S_0) \to 1$ almost surely as $n \to \infty$ for any t_l. Since there are finite number of such items in $a(n, j)$, it must converge to 1 almost surely as n goes to infinity.

As a result, $[S_i^*(t_1, n), \ldots, S_i^*(t_m, n)]$ and $[\hat{S}_i(t_1), \ldots, \hat{S}_i(t_m)]$ for any $i = 1, \ldots, n$ can be made arbitrarily close, almost surely, for a large enough n. For the contingent payoff, however, the difference between the pre- and post-adjustment values depends on $[\hat{S}_i(t_1), \ldots, \hat{S}_i(t_m)]$ because the EMS adjustment is multiplicative. Convergence in the contingent payoff value therefore needs Lipschitz continuity. Since $f(\cdot)$ is Lipschitz continuous by assumption, the following result is immediate:

$$\begin{aligned}
&|f[S_i^*(t_1, n), \ldots, S_i^*(t_m, n)] - f[\hat{S}_i(t_1), \ldots, \hat{S}_i(t_m)]| \\
&\le c\|[S_i^*(t_1, n), \ldots, S_i^*(t_m, n)] - [\hat{S}_i(t_1), \ldots, \hat{S}_i(t_m)]\| \\
&\le c \sum_{j=1}^{m} |S_i^*(t_j, n) - \hat{S}_i(t_j)| \\
&= c \sum_{j=1}^{m} |\hat{S}_i(t_j)| \times |a(n, j) - 1|.
\end{aligned}$$

This in turn implies that

$$\begin{aligned}
&\frac{1}{n} \sum_{i=1}^{n} |f[S_i^*(t_1, n), \ldots, S_i^*(t_m, n)] - f[\hat{S}_i(t_1), \ldots, \hat{S}_i(t_m)]| \\
&\le \frac{c}{n} \sum_{i=1}^{n} \sum_{j=1}^{m} |\hat{S}_i(t_j)| \times |a(n, j) - 1| \\
&= c \sum_{j=1}^{m} |a(n, j) - 1| \left(\frac{1}{n} \sum_{i=1}^{n} |\hat{S}_i(t_j)| \right) \to 0 \\
&\quad \text{almost surely.}
\end{aligned}$$

Up to this point, we have proved the EMS estimate is asymptotically equivalent to the standard Monte Carlo estimate. It remains to show that the standard Monte Carlo estimate converges to the theoretical value under the assumption of Lipschitz continuous payoff function. Use $[S_0(t_1), \ldots, S_0(t_m)]$ to denote an arbitrary nonrandom sample path. We use Lipschitz continuity to obtain

$$\begin{aligned}
&|f[S(t_1), \ldots, S(t_m)]| - |f[S_0(t_1), \ldots, S_0(t_m)]| \\
&\le |f[S(t_1), \ldots, S(t_m)] - f[S_0(t_1), \ldots, S_0(t_m)]| \\
&\le c \sum_{j=1}^{m} |S(t_j) - S_0(t_j)| \\
&\le c \sum_{j=1}^{m} |S(t_j)| + c \sum_{j=1}^{m} |S_0(t_j)|.
\end{aligned}$$

This implies that

$$\begin{aligned}
&E^Q\{|f[S(t_1), \ldots, S(t_m)]| \mid \mathcal{F}_0\} \\
&\le c \sum_{j=1}^{m} E^Q\{|S(t_j)| \mid \mathcal{F}_0\} \\
&\quad + c \sum_{j=1}^{m} |S_0(t_j)| + |f[S_0(t_1), \ldots, S_0(t_m)]|.
\end{aligned}$$

Since $E^Q\{|S(t_j)| \mid \mathcal{F}_0\}$ for $j = 1, \ldots, m$ are finite by the martingale assumption, $E^Q\{|f[S(t_1), \ldots, S(t_m)]| \mid \mathcal{F}_0\}$ must also be finite. Due to random sampling, $f[\hat{S}_i(t_1), \ldots, \hat{S}_i(t_m)]$ for $i = 1, 2, \cdots$ form a *i.i.d.* sequence. Together, we can apply ergodic theorem to conclude that

$$\begin{aligned}
&\frac{1}{n} \sum_{i=1}^{n} f[\hat{S}_i(t_1), \ldots, \hat{S}_i(t_m)] \to \\
&\quad E^Q\{|f[S(t_1), \ldots, S(t_m)]| \mid \mathcal{F}_0\} \quad \text{almost surely.}
\end{aligned}$$

The desired convergence result is thus established. □

Proof of Proposition 2

Let A_l, $l = 1, \ldots, k$ be the elements of the partition defining the generic Lipschitz continuity. We now define a boundary set as follows: $W = \cup_{l=1}^{k} [C(A_l) - I(A_l)]$ where $C(\cdot)$ and $I(\cdot)$ denote closure and interior, respectively. Note that W is a closed set. Corresponding to the ith sample path $[\hat{S}_i(t_1), \ldots, \hat{S}_i(t_m)]$, we define x_i to be an element of W that $\|x_i - [\hat{S}_i(t_1), \ldots, \hat{S}_i(t_m)]\| \le \|y - [\hat{S}_i(t_1), \ldots, \hat{S}_i(t_m)]\|$ for any $y \in W$. Let $B(x_i, \delta)$ denote an open ball around x_i. For any $u > 0$ and $\delta > 0$, there exists an N such that for any $n \ge N$, $\|[a(n, 1) - 1, a(n, 2) - 1, \ldots, a(n, m) - 1]\| < (\delta/u)$ almost surely because $a(n, j)$, as defined earlier in the proof for Proposition 1, converges to 1 almost surely.

We first show that if $[\hat{S}_i(t_1), \ldots, \hat{S}_i(t_m)] \notin B(x_i, \delta)$, $\|[\hat{S}_i(t_1), \ldots, \hat{S}_i(t_m)]\| \le u$ and $[\hat{S}_i(t_1), \ldots, \hat{S}_i(t_m)] \in A_l$, it must be that $[S_i^*(t_1, n), \ldots, S_i^*(t_m, n)] \in A_l$ for $n \ge N$ as above. Suppose the contrary. There must exist a $\lambda \in [0, 1]$ such that $\bar{y} \equiv \lambda[S_i^*(t_1, n), \ldots, S_i^*(t_m, n)] + (1 - \lambda)[\hat{S}_i(t_1), \ldots, \hat{S}_i(t_m)] \in W$. By the definition of x_i and the fact that $[\hat{S}_i(t_1), \ldots, \hat{S}_i(t_m)] \notin B(x_i, \delta)$, we have $\|\bar{y} - [\hat{S}_i(t_1), \ldots, \hat{S}_i(t_m)]\| \ge \|x_i - [\hat{S}_i(t_1), \ldots, \hat{S}_i(t_m)]\| \ge \delta$. Moreover, for $n \ge N$,

$$\begin{aligned}
&\|\bar{y} - [\hat{S}_i(t_1), \ldots, \hat{S}_i(t_m)]\| \\
&= \|\lambda\{[a(n, 1) - 1]\hat{S}_i(t_1), [a(n, 2) - 1]\hat{S}_i(t_2), \ldots, [a(n, m) - 1]\hat{S}_i(t_m)\}\| \\
&\le \lambda\|[a(n, 1) - 1, a(n, 2) - 1, \ldots, a(n, m) - 1]\| \\
&\quad \times \|[\hat{S}_i(t_1), \ldots, \hat{S}_i(t_m)]\| \\
&\le \lambda u\|[a(n, 1) - 1, a(n, 2) - 1, \ldots, a(n, m) - 1]\| \\
&< \lambda\delta \\
&\le \delta,
\end{aligned}$$

which is a contradiction.

To simplify exposition, we may use $\hat{S}_i$ and $S_i^*(n)$ to denote $[\hat{S}_i(t_1), \ldots, \hat{S}_i(t_m)]$ and $[S_i^*(t_1, n), \ldots, S_i^*(t_m, n)]$, respectively. For $n \ge N$,

$$\frac{1}{n}\sum_{i=1}^{n} |f(S_i^*(n)) - f(\hat{S}_i)|$$
$$= \frac{1}{n}\sum_{i=1}^{n} \chi_{\{\|\hat{S}_i\|\le u\}}\chi_{\{\hat{S}_i\notin B(x_i,\delta)\}} |f(S_i^*(n)) - f(\hat{S}_i)|$$
$$+ \frac{1}{n}\sum_{i=1}^{n} \chi_{\{\|\hat{S}_i\|\le u\}}\chi_{\{\hat{S}_i\in B(x_i,\delta)\}} |f(S_i^*(n)) - f(\hat{S}_i)|$$
$$+ \frac{1}{n}\sum_{i=1}^{n} \chi_{\{\|\hat{S}_i\|> u\}} |f(S_i^*(n)) - f(\hat{S}_i)|,$$

where $\chi_{\{\cdot\}}$ denotes the indicator function.

The first term becomes

$$\frac{1}{n}\sum_{i=1}^{n} \chi_{\{\|\hat{S}_i\|\le u\}}\chi_{\{\hat{S}_i\notin B(x_i,\delta)\}} |f(S_i^*(n)) - f(\hat{S}_i)|$$
$$\le \frac{c}{n}\sum_{i=1}^{n} \chi_{\{\|\hat{S}_i\|\le u\}}\chi_{\{\hat{S}_i\notin B(x_i,\delta)\}} \|S_i^*(n) - \hat{S}_i\|$$
$$\le \frac{c}{n}\sum_{i=1}^{n} \|S_i^*(n) - \hat{S}_i\|$$
$$\le \frac{c}{n}\sum_{i=1}^{n}\sum_{j=1}^{m} |\hat{S}_i(t_j)| \times |a(n,j) - 1|$$
$$= c\sum_{j=1}^{m} |a(n,j) - 1| \left(\frac{1}{n}\sum_{i=1}^{n} |\hat{S}_i(t_j)|\right) \to 0$$
$$\text{almost surely as } n \to \infty.$$

The first inequality follows from Lipschitz continuity for points in the same element of the partition. The derivation also utilizes ergodic theorem and the fact that $a(n, j)$ converges to 1 almost surely.

The second term becomes

$$\frac{1}{n}\sum_{i=1}^{n} \chi_{\{\|\hat{S}_i\|\le u\}}\chi_{\{\hat{S}_i\in B(x_i,\delta)\}} |f(S_i^*(n)) - f(\hat{S}_i)|$$
$$\le \frac{b}{n}\sum_{i=1}^{n} \chi_{\{\|\hat{S}_i\|\le u\}}\chi_{\{\hat{S}_i\in B(x_i,\delta)\}} (2 + \|S_i^*(n)\| + \|\hat{S}_i\|)$$
$$\le \frac{b}{n}\sum_{i=1}^{n} \chi_{\{\|\hat{S}_i\|\le u\}}\chi_{\{\hat{S}_i\in B(x_i,\delta)\}} \left[2 + \left(1 + \frac{\delta}{u}\right)\|\hat{S}_i\|\right]$$
$$\le \frac{b(2+\delta+u)}{n}\sum_{i=1}^{n} \chi_{\{\hat{S}_i\in B(x_i,\delta)\}} \to b(2+\delta+u)\Pr\{\hat{S}_i \in B(x_i,\delta)\}$$
$$\text{almost surely as } n \to \infty$$
$$\le b(2+\delta+u)K\delta^m \text{ where } K \text{ is some finite positive constant.}$$

The first inequality is due to the generic Lipschitz continuity whereas the second one is due to $n \ge N$. The restriction $\|\hat{S}_i\| \le u$ ensures the third inequality. The convergence result is due to ergodic theorem. The last inequality is a result of the bounded multivariate density function.

The third term becomes

$$\frac{1}{n}\sum_{i=1}^{n} \chi_{\{\|\hat{S}_i\|> u\}} |f(S_i^*(n)) - f(\hat{S}_i)|$$
$$\le \frac{b}{n}\sum_{i=1}^{n} \chi_{\{\|\hat{S}_i\|> u\}} (2 + \|S_i^*(n)\| + \|\hat{S}_i\|)$$
$$\le \frac{b}{n}\sum_{i=1}^{n} \chi_{\{\|\hat{S}_i\|> u\}} \left[2 + \left(1 + \frac{\delta}{u}\right)\|\hat{S}_i\|\right] \to$$
$$2b\Pr\{\|\hat{S}_i\| > u\} + b\left(1 + \frac{\delta}{u}\right)E^Q\{\chi_{\{\|\hat{S}_i\|>u\}}\|\hat{S}_i\| \mid \mathcal{F}_0\}.$$

The first inequality is again due to the generic Lipschitz continuity whereas the second one due to $n \ge N$. The convergence result is due to ergodic theorem.

Since u and δ are arbitrary, we can set $\delta = 1/u^2$. For any u, there is a corresponding $N(u)$ such that for any $n \ge N(u)$, the above results pertaining to the three terms are true. Let $u \to \infty$, the second term converges to 0 because $u\delta^m = u^{1-2m} \to 0$. The third term goes to 0 when $u \to \infty$. We thus have

$$\frac{1}{n}\sum_{i=1}^{n} |f(S_i^*(n)) - f(\hat{S}_i)| \to 0 \text{ almost surely as } n \to \infty.$$

It remains to show that $(1/n)\sum_{i=1}^{n} f[\hat{S}_i(t_1), \ldots, \hat{S}_i(t_m)]$ converges to $E^Q\{|f([S(t_1), \ldots, S(t_m)]| \mid \mathcal{F}_0\}$ almost surely. Under the assumption of the generic Lipschitz condition, we have

$$|f[S(t_1), \ldots, S(t_m)]| \le b(1 + \|[S(t_1), \ldots, S(t_m)]\|),$$

and thus,

$$E^Q\{|f[S(t_1), \ldots, S(t_m)]| \mid \mathcal{F}_0\}$$
$$\le b(1 + E^Q\{\|[S(t_1), \ldots, S(t_m)]\| \mid \mathcal{F}_0\}).$$

Since $E^Q\{\|[S(t_1), \ldots, S(t_m)]\| \mid \mathcal{F}_0\}$ is finite by the martingale assumption, $E^Q\{|f[S(t_1), \ldots, S(t_m)]| \mid \mathcal{F}_0\}$ must also be finite. Due to random sampling, $f[\hat{S}_i(t_1), \ldots, \hat{S}_i(t_m)]$ for $i = 1, 2, \cdots$ form an i.i.d. sequence. We can now apply ergodic theorem to conclude that

$$\frac{1}{n}\sum_{i=1}^{n} f[\hat{S}_i(t_1), \ldots, \hat{S}_i(t_m)] \to$$
$$E^Q\{|f[S(t_1), \ldots, S(t_m)]| \mid \mathcal{F}_0\} \quad \text{almost surely.}$$

The desired result is therefore established. □

References

Barraquand, J. 1995. Numerical valuation of high dimensional multivariate european securities. *Management Sci*. **41** 1882–1891.

Black, F., M. Scholes 1973. The pricing of options and corporate liabilities. *J. Political Econom*. **81** 637–659.

Boyle, P. 1977. Options: a monte carlo approach. *J. Financial Econom.* **4** 323–338.

——, P., M. Broadie, P. Glasserman 1997. Monte carlo methods for security pricing. *J. Econom. Dynam. Control.* **21** 1263–1321.

Broadie, M., J. Detemple 1996. American option valuation: new bounds, approximations, and a comparison of existing methods. *Rev. Financial Stud.* **9** 1211–1250.

Cox, J., S. Ross 1976. The valuation of options for alternative stochastic processes. *J. Financial Econ.* **3** 145–166.

Duan, J. 1995. The GARCH option pricing model. *Math. Finance.* **5** 13–32.

Harrison, J., D. Kreps 1979. Martingales and arbitrage in multiperiod securities markets. *J. Econom. Theory.* **20** 381–408.

——, S. Pliska 1981. Martingales and stochastic integrals in the theory of continuous trading. *Stochastic Process. Appl.* **11** 215–260.

Hull, J., A. White 1987. The pricing of options on assets with stochastic volatility. *J. Finance.* **42** 281–299.

Kemna, A., A. Vorst 1990. A pricing method for options based on average asset values. *J. Banking and Finance.* **14** 113–129.

Lucas, R. 1978. Asset prices in an exchange economy. *Econometrica.* **46** 1429–1445.

Merton, R. 1973. Theory of rational option pricing. *Bell J. Econom. and Management Sci.* **4** 141–183.

Ritchken, P., L. Sankarasubramanian, A. Vijh 1993. The valuation of path dependent contracts on the average. *Management Sci.* **39** 1202–1213.

Turnbull, S., L. Wakeman 1991. A quick algorithm for pricing european average options. *J. Financial and Quantitative Anal.* **26** 377–389.

Accepted by Robert Heinkel; received January 22, 1996. This paper has been with the authors $3\frac{1}{2}$ *months for 2 revisions.*