DEF 24, 107 – 126 (2001)

Decisions in Economics and Finance

Efficient Monte Carlo pricing of European options using mean value control variates

P. Pellizzari[*]

Department of Applied Mathematics, University of Venice
e-mail: `paolop@unive.it`

Received: 14 July 2000 / Accepted: 18 February 2001

Abstract. We describe in this paper a variance reduction method based on control variates. The technique uses the fact that, if all stochastic assets but one are replaced in the payoff function by their mean, the resulting integral can most often be evaluated in closed form. We exploit this idea by applying the univariate payoff as control variate and develop a general Monte Carlo procedure, called Mean Monte Carlo (MMC). The method is then tested on a variety of multifactor options and compared to other Monte Carlo approaches or numerical techniques. The method is of easy and broad applicability and gives good results especially for low to medium dimension and in high volatility environments.

Mathematics Subject Classification (2000): 65C05, 91B28

Journal of Economic Literature Classification: C15, G12

1. Introduction

The valuation of complex derivatives has been the subject of a huge research effort in recent years. Problems like American exercise, path-dependency, and non-standard stochastic processes challenge the well-known Black–Scholes solution in non-trivial ways and numerical techniques are still actively investigated to find approximate risk-neutral valuations. Even taking

---

[*] We thank the participants at the symposium “Mathematical methods in applied finance”, the SIAM 99 annual meeting, S. Hodges for useful discussions, two anonymous referees and the editor for constructive criticism and useful bibliographic references. This work was partially supported by MURST.

into account only European options, as it will be the case in this paper, we find that the presence of many underlying sources of risk does not allow straightforward application of analytical valuation formulae. In Cox and Ross (1976) it is noted that, if a riskless hedge can be formed, the price can be computed as the discounted risk-neutral expectation of the payoff of the claim. In the case of European derivatives, this leads to the computation of a multidimensional integral for which most often there is no known result in closed form. Monte Carlo (MC) methods can be used successfully in this framework and early pricing applications date back to Boyle (1977) and Kemna and Vorst (1990). The main advantages of MC over other methods are simplicity and broad applicability. In particular, while it is hard to solve the partial differential equation for the price (see Wilmott et al. (1995) for an introduction) or to use tree based methods (Cox et al. (1979)) when there are more than three underlying assets, MC can still easily be applied with little difficulty. On the other hand, the error of MC methods converges to zero as the reciprocal of the square root of the number of simulations. Hence, virtually every other method is faster, often by orders of magnitude. This fact has been acknowledged from the beginning of the MC applications to finance as both Boyle and Kemna and Vorst papers contain accelerating tricks to speed up the computations.

In this paper, we propose a control variates approach to reduce the variance of estimated prices of multidimensional claims. The key idea is to exploit the closed valuation formula usually known in the univariate case to reduce the variance of the simulated multivariate payoff. This is done by replacing all the arguments but one by their means; hence we name our method *Mean Monte Carlo* (MMC). A similar but simpler technique is described in Pellizzari (2001), where only portfolio options are examined and no optimization is performed to reduce variance.

To our knowledge, little work has been done to develop general control variates for MC pricing; rather each solution was somehow chosen *ad hoc* as like in Boyle (1977) or Kemna and Vorst (1990). An influential paper that deals with European multivariate option pricing is Barraquand (1995). The proposed method is a resampling technique, named “quadratic resampling” (QR), that can in principle be applied to every simulation problem. The random sample is modified in such a way as to have first and second moments exactly equal to the theoretical values, thus providing “second order” approximation to the ideal sample. Similar moment matching ideas are also presented in Boyle et al. (1997). Our method differs from the above contributions as we do not propose a new sampling procedure but rather define a way to improve existing (stochastic or deterministic) sampling schemes. The only requirement for the use of MMC is easy analytical evaluation of the univariate restriction as is available for many common options. Note that, even if numerical approximations should be used to solve a univariate integral, there are extremely accurate methods (e.g., Gaussian quadrature) for univariate problems. See Krommer and Ueberhuber (1998) for a comprehensive account.

A different idea is exploited in Milevsky and Posner (1998a), where a valuation procedure for portfolio and Asian options based on an approximation of the payoff density is presented. It is argued that the reciprocal gamma distribution can approximate the sum of lognormals, providing a closed-form expression for the price. The authors claim that this approximation is more accurate than the lognormal one and test their procedure against MC results.

The rest of the paper is organized as follows. In Section 2 we present the MMC method, by means of a simple example of an option written on three correlated assets. Theoretical discussion is provided linking the required positive correlation between control and payoff to stochastic orders concepts. Section 3 contains pricing applications to a wide variety of multivariate options. We begin by evaluating an exchange option on two assets as this is one of the few cases for which an analytical pricing formula exists and which can be used as a benchmark. We then contrast our results with Barraquand's, showing that there is no dominant method and then price some portfolio options. We compare MMC to the Milevsky and Posner methodology on a 7-dimensional basket option, related to index-linked guaranteed investment certificates. We then perform a simulation study pricing a sample of 1000 low (2- to 6-) and high (11- to 40-) dimensional portfolio options. In these and subsequent applications, as outlined in Broadie and Detemple (1996), the quality of our results is assessed by running MMC on a set of options whose parameters are randomly selected. In Section 3.4 we study "multiple" options and propose a rule of thumb to select two effective control variates. Finally we give some concluding remarks in Section 4.

2. Mean variance reduction

Let $S_1, \dots, S_n$ be stocks whose risk-neutral processes follow the stochastic differential equation

$$dS_i = (r - d_i)S_i dt + \sigma_i S_i dZ_i,$$

where r is the instantaneous riskless rate, d_i is the continuously paid dividend yield, σ_i is the volatility and Z_i is a standard Brownian motion.

Assume that we want to price at time $t = 0$ a European-like asset that pays the sum

$$C_T = f(S_{1T}, \dots, S_{nT}), \tag{1}$$

at time T, where S_{it}, $0 \leq t \leq T$, denotes the value of i-th asset at time t.

The price $P_T = \exp(-rT)E\left[C_T\right]$ of the claim (1) can be estimated by generating many realizations of $\left\{S_{1T}^{(j)}, \ldots, S_{nT}^{(j)}\right\}$, $j = 1, \ldots, N$, and discounting the sample mean of the resulting $C_T^{(j)} = f\left(S_{1T}^{(j)}, \ldots, S_{nT}^{(j)}\right)$, $j = 1, \ldots, N$, to get

$$\hat{C}_T = e^{-rT} \frac{1}{N} \sum_j C_T^{(j)}.$$

The computational cost is increasing linearly in the dimension and can be substantial, as many vector random variables are to be generated from a multivariate distribution. Even more importantly, the standard deviation of the estimated price is driven by the $O(1/\sqrt{N})$ law of large numbers and hence a huge N might be required to achieve satisfactory precision.

We describe a simple implementation of a variance reduction scheme based on control variates. A candidate control variate is a random variable positively correlated with C_T and such that its expected value can be readily computed (see Rubinstein (1981) or Ross (1997) for an introduction and Schmeiser et al. (2001, 2002) for recent generalizations). Consider the following control variates $M_T(i), i = 1, \ldots, n$:

$$M_T(i) = f(E\left[S_{1T}\right], \ldots, E\left[S_{i-1,T}\right], S_{iT}, E\left[S_{i+1,T}\right], \ldots, E\left[S_{nT}\right]). \quad (2)$$

The random variable $M_T(i)$ is obtained from (1) by replacing S_{jT} with its unconditional mean $E[S_{jT}]$ if $i \neq j$. It is obvious that $M_T(i)$ is in general correlated with C_T and its mean $E[M_T(i)]$ can be easily evaluated in many important cases being the expectation the result of a univariate integration.

The control variates (2) allow us to obtain a set of Monte Carlo estimates $\hat{C}_T(i)$ of the unknown price P_T:

$$\hat{C}_T(i) = e^{-rT} \frac{1}{N} \sum_{j=1}^{N} \left[C_T^{(j)} - M_T^{(j)}(i) + E\left[M_T(i)\right]\right], \quad i = 1, \ldots, n,$$

where $C_T^{(j)}$ and $M_T^{(j)}(i)$ denote the j-th realizations of the payoff C_T and control variate $M_T(i)$.

*Example 2.1 (*Pellizzari (2001)*).* Assume that we want to price portfolio options on three assets, with payoff

$$f(S_{1T}, S_{2T}, S_{3T}) = \max\{0, q_1 S_{1T} + q_2 S_{2T} + q_3 S_{3T} - k\},$$

where the q_i's are the positive quantities held in each asset. The three control variates $M_T(i), i = 1, 2, 3$, are

$$\begin{aligned} M_T(1) &= \max\left\{0, q_1 S_{1T} + q_2 E\left[S_{2T}\right] + q_3 E\left[S_{3T}\right] - k\right\}, \\ M_T(2) &= \max\{0, q_1 E\left[S_{1T}\right] + q_2 S_{2T} + q_3 E\left[S_{3T}\right] - k\}, \\ M_T(3) &= \max\{0, q_1 E\left[S_{1T}\right] + q_2 E\left[S_{2T}\right] + q_3 S_{3T} - k\}. \end{aligned} \quad (3)$$

The previous assumptions yield

$$E\left[S_{iT}\right] = S_{i0} \exp\left(\left(r - d_i\right) T\right), \quad i = 1, \ldots, 3,$$

where the S_{i0} are the initial values of i-th assets and the means of the control variates $M_T(i), i = 1, 2, 3$, can be readily evaluated. Note, for example, that the right-hand side of (3) can be rewritten as

$$q_3 \max\left\{0, S_{3T} - \frac{k - q_1 E\left[S_{1T}\right] - q_2 E\left[S_{2T}\right]}{q_3}\right\},$$

which can be recognized easily as the payoff of a European call option on the asset S_3 and strike price $K_3 = (k - q_1 E\left[S_{1T}\right] - q_2 E\left[S_{2T}\right])/q_3$. Hence the expected value can be derived by the Black–Scholes formula.

A straightforward generalization of the one control variate alone method is based on the use of many $M_T(i)$ at the same time, which provide the prices

$$\hat{C}_T(A, \mathbf{b}) = e^{-rT} \frac{1}{N} \sum_{j=1}^{N} \left[f(S_{1T}^{(j)}, \ldots, S_{nT}^{(j)}) \right.$$
$$\left. + \sum_{i \in A} b_i \left(M_T^{(j)}(i) - E\left[M_T(i)\right] \right) \right],$$

depending on the set $A \subseteq \{1, 2, \ldots, n\}$ of indices used in variance reduction and on the vector $\mathbf{b}$. The best price should be obtained by selecting the set A and coefficient vector $\mathbf{b}$ so that

$$\min_{A \subseteq \{1,2,\ldots,n\}, \mathbf{b}} \operatorname{Var}(\hat{C}_T(A, \mathbf{b})) \tag{4}$$

is attained. This minimization problem is interesting because both subset selection (on A) and numerical optimization (on $\mathbf{b}$) are involved. There is a standard way to determine $\mathbf{b}$ by regression for given A (see Boyle et al. (1997)), but the joint problem of optimal subset selection is harder. Hence, in the remainder of the paper we set $A = \{1, \ldots, n\}$ unless otherwise stated and solve problem (4) with respect to $\mathbf{b}$ whenever control variates are used. The resulting Monte Carlo method will be denoted in what follows by *Mean Monte Carlo* (MMC). Any control variate technique with optimized coefficients reduces the variance of the vanilla MC estimate, provided that the control and the payoff are correlated, which is proved in an elementary way in Pellizzari (2001). It turns out that this result is a particular case of Proposition 9.C.3 in Shaked and Shantikumar (1994): in fact if the vector of returns is weakly conditionally increasing in sequence (WCIS) then it is associated and this implies the desired positive correlation.

Finally, note that this picture does not take into account finer details related to computation time: in the end, the objective of any variance reduction technique is to provide an estimate within given accuracy in the shortest time. It is common to evaluate performance by comparing standard deviations (or variances), somehow assuming that different methods all run in the same time and we will not depart totally from this custom. This is perhaps a sensible approximation, but it is wrong. To exemplify this in our setting: it might well be the case that the use of two $M_T(i)$'s produces (slightly) smaller variance than one alone, but at the price of longer computation time that could question (or even reverse) the apparent superiority of one method.

3. Applications

The present section contains examples and applications of the control variates previously defined. The method is comprehensively applied to different multivariate option pricing problems. We first check the results obtained by MMC against a closed evaluation formula, using the 2-dimensional exchange option of Margrabe (1978). As far as we know, this is one of the few multivariate cases (another being the max option studied in Stulz (1982)) for which an analytical pricing formula is known. We then compare MMC to the pricing methodology of Barraquand (1995) on some 5-, 10-, and 100-dimensional options on the maximum of several assets. In what follows we turn to some applications to portfolio options examining a claim written on 7 assets, whose price is estimated by approximating the risk-neutral density by reciprocal gamma and lognormal distributions (Milevsky and Posner (1998)). Moreover, we perform a simulation study pricing up to 6 and up to 40 asset portfolios. In the first case, a comparison is also made against quasi-MC deterministic integration methods (based on Sobol sequences). The latter is an application to high dimensional problems that enables us to show the strengths and limits of our method. In Section 3.4, we price some "multiple" options, having payoff function $\max(0, S_{1T} - k_1, \ldots, S_{nT} - k_n)$, and present an effective rule of thumb selection of the set A of control variates. In all cases the optimization problem (4) is approximately solved by regressing 1024 multidimensional control variates on the payoff values.

3.1. *Exchange option*

We consider an exchange option (Margrabe (1978)) on two assets whose payoff is

$$f(S_{1T}, S_{2T}) = \max\{0, S_{2T} - S_{1T}\},$$

and for which the following analytic pricing formula is available:

$$P_T = S_{20}e^{-d_2T}N(p) - S_{10}e^{-d_1T}N(p - \Sigma\sqrt{T}),$$

where $N(\cdot)$ is the cumulative standard normal distribution and

$$p = \frac{\log(\frac{S_{20}e^{-d_2T}}{S_{10}e^{-d_1T}})}{\Sigma\sqrt{T}} + \frac{1}{2}\Sigma\sqrt{T}, \quad \Sigma^2 = \sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2.$$

Table 1. Estimated prices and relative standard deviations in brackets for various MC methods

N	MC	SMC	AMC	MMC
1024	15.20 (0.68)	16.00 (0.70)	16.30 (0.35)	16.11 (0.15)
2048	15.57 (0.49)	16.06 (0.50)	16.42 (0.25)	16.12 (0.11)
4096	15.80 (0.35)	16.08 (0.35)	15.98 (0.17)	16.09 (0.08)
8192	16.17 (0.25)	16.06 (0.25)	16.25 (0.13)	16.12 (0.05)
True	16.06			

Setting $S_{10} = S_{20} = 100$, $r = \log(1.1), d_1 = d_2 = \log(1.05)$, $\sigma_1 = 0.3, \sigma_2 = 0.2$, $\rho = -0.5$ and $T = 0.95$, we get the price 16.0606. Table 1 shows some estimates obtained by plain MC, deterministic Sobol Monte Carlo (SMC), standard antithetic MC (AMC) and MMC with standard deviations in brackets for different sample sizes N. AMC is described in Hammersley and Handscombe (1967), Rubinstein (1981) and other textbooks on Monte Carlo methods and is implemented most commonly in an MC simulation. A brief description of SMC will be deferred until Section 3.3 on portfolio options. MMC produces an error 4 to 5 times smaller than MC and SMC methods, while it is “just” twice as accurate as AMC. Note also that the estimates produced by SMC are very close to the true price and we will argue in the sequel that standard deviation is not entirely informative as an error measure in this case.

3.2. Quadratic resampling

Barraquand (1995) describes a multidimensional integration method based on a resampling technique that matches the first two moments of the empirical distribution of sampled values with the theoretical ones. In this respect, the method is a generalization in multiple dimensions of moment matching algorithms presented in Boyle et al. (1997); in particular, QR resembles moment matching of second order on the values $S_{1T}, \ldots, S_{nT}$. QR computes exactly multidimensional integrals involving powers of degree up to 2 in the variables. This is quite intuitive, due to the above two-moments property. Further details can also be found in Barraquand (1993), an extended and more user friendly version of Barraquand (1995). The method is tested in the following framework: let the covariance matrix of the multivariate normal vector $X = (x_1, \dots, x_n)^T$ be $H = (h_{ij}), i, j = 1, \dots, n$, where

$$h_{ii} = \sigma^2,\ i = 1, \dots, n, \quad \text{and} \quad h_{ij} = \rho\sigma^2,\ i \neq j. \tag{5}$$

In order to ensure the positive semi-definiteness of H, we must have $-1/(n-1) \leq \rho \leq 1$. Let V be the Cholesky decomposition of H, i.e.,

$$H = VV^T.$$

With the function f defined as

$$f(x_1, \dots, x_n) = \max_{i=1}^{n} \exp(x_i),$$

our aim is the calculation of the multidimensional integral

$$I(f) = E(f(X)) = \int_{\mathbf{R}^n} f(X)p_X(X)dX = \int_{\mathbf{R}^n} f(VZ)p_Z(Z)dZ,$$

where p_X and p_Z are respectively the densities of the random multinormal vector X and the multinormal standard vector Z. More explicitly, the above integral can be written as

$$I(n, \sigma, \rho) = \int_{\mathbf{R}^n} \max_{i=1}^{n} \exp\left(\sum_{j=1}^{n} v_{ij} z_j \frac{\exp\left(-\sum_{k=1}^{n} z_k^2/2\right)}{\sqrt{2\pi}^{n}}\right) dz_1 \dots dz_n$$

This integral can be interpreted, up to multiplication by $\exp(\sigma^2/2)$, as the price of an option on the maximum of n assets whose log-returns are X, time to expiration $T = 1$ and riskfree rate $r = 0$.

The results listed in Barraquand (1995) are hard to evaluate and replicate, as they are obtained using three variance reducing methods at the same time, namely, antithetic variates, quadratic resampling and additive importance sampling. As it is not trivial to attribute to one single method all the merits of a good performance, we use the results reported in Barraquand (1993) under the denomination "Method 3", where just two techniques (antithetic variates and quadratic resampling) are simultaneously applied.

Hence, Table 2 contrasts the results obtained by antithetic QR and antithetic MMC, for $n = 5, 10, 100$, $\sigma = 0.1, 1.0$, number of simulations $N = 4000$ and various values of correlations ρ. We also provide AMC and MC figures as benchmarks.

Overall, MMC is superior to QR in 15 instances out of 22 if the standard deviation (SD) is used to measure performance. As we do not know the running times of the QR experiments, we are unable to perform more detailed

Table 2. Comparison of Barraquand results and MMC prices, standard deviations in brackets

n	σ	ρ	QR+AMC	MMC+AMC	AMC	MC
5	0.1	$-\frac{1}{4}$	1.141 (0.00027)	1.142 (0.00035)	1.142 (0.00090)	1.140 (0.00101)
		0.0	1.126 (0.00024)	1.127 (0.00044)	1.126 (0.00079)	1.126 (0.00123)
		0.5	1.090 (0.00017)	1.089 (0.00049)	1.089 (0.00053)	1.092 (0.00148)
		1.0	1.005 (4e-7)	1.005 (0.00004)	1.005 (0.00011)	1.004 (0.00158)
	1.0	$-\frac{1}{4}$	4.39 (0.0230)	4.39 (0.0079)	4.46 (0.0465)	4.36 (0.0532)
		0.0	4.07 (0.0255)	4.07 (0.0111)	4.11 (0.0389)	4.11 (0.0583)
		0.5	3.29 (0.0205)	3.28 (0.0136)	3.31 (0.0339)	3.37 (0.0589)
		1.0	1.65 (0.067)	1.65 (0.001)	1.65 (0.019)	1.62 (0.033)
10	0.1	$-\frac{1}{9}$	1.178 (0.00050)	1.177 (0.00047)	1.179 (0.00080)	1.178 (0.00099)
		0.0	1.168 (0.00047)	1.168 (0.00047)	1.167 (0.00074)	1.169 (0.00112)
		0.5	1.119 (0.00032)	1.118 (0.00048)	1.119 (0.00050)	1.118 (0.00148)
		1.0	1.005 (4.8e-7)	1.005 (0.00003)	1.005 (0.00011)	1.005 (0.00148)
	1.0	$-\frac{1}{9}$	5.91 (0.0423)	5.91 (0.0166)	5.96 (0.0504)	5.93 (0.0698)
		0.0	5.62 (0.0424)	5.61 (0.0180)	5.54 (0.0459)	5.70 (0.0713)
		0.5	4.18 (0.0287)	4.17 (0.0187)	4.22 (0.0371)	4.19 (0.0666)
		1.0	1.65 (0.0083)	1.65 (0.0010)	1.64 (0.0164)	1.64 (0.0360)
100	0.1	0.0	1.286 (0.00070)	1.287 (0.00060)	1.285 (0.00062)	1.285 (0.00089)
		0.5	1.198 (0.00046)	1.197 (0.00043)	1.198 (0.00042)	1.197 (0.00148)
		1.0	1.005 (5.1e-7)	1.005 (0.00009)	1.005 (0.00011)	1.005 (0.00158)
	1.0	0.0	13.58 (0.0994)	13.48 (0.0756)	13.47 (0.0812)	13.50 (0.1114)
		0.5	7.94 (0.0486)	7.84 (0.0449)	8.00 (0.0620)	8.03 (0.1211)
		1.0	1.65 (0.0087)	1.65 (0.0010)	1.64 (0.0194)	1.64 (0.0329)

comparisons. MMC appears to outperform QR in 11 cases out of 11 if high volatility is present.

We conclude with some considerations on the examples chosen by Barraquand. From a financial point of view $I(n, \sigma, \rho)$ is, up to a constant, the price of an option on the maximum among n correlated assets. However, due to no strike price, extreme correlation and volatilities, this interpretation provides terribly in the money options. Just think that some options are worth around 13 against unitary initial values of assets. Besides being of questionable financial interest, we recall that such extreme circumstances are very rare in practical option pricing.

3.3. Portfolio options

3.3.1. Milevsky and Posner approximation An approximation of the payoff risk-neutral density can be used to provide estimates of true prices, if this enables us to obtain analytical results for the integral that defines the expectation. This approach is taken in Milevsky and Posner (1998a) to price portfolio and Asian options. It is argued that the distribution of sums of lognormals can be approximated by a reciprocal gamma distribution. This is motivated by the fact that the density of the sum of correlated lognormals tends to this distribution for some “decaying” correlation structure. We do not provide full details, but refer the reader to the paper, which also describes an application to a 7-dimensional portfolio option embedded in several financial products, called Index-Linked Guaranteed Investment Certificates (ILGIC). In brief, the final return of an ILGIC is related to the performance of a basket of some underlying indices. Consider a call option that pays the max of a weighted average of normalized stock indices listed in Table 3.

Table 3. Stock indices included in an ILGIC, with their weights, volatilities and dividend yields

Country	Index	Weight (%)	Ann. Vol. (%)	Div. Yield (%)
Canada	TSE 100	10	11.55	1.69
France	CAC 40	15	20.68	2.39
Germany	DAX	15	14.53	1.36
Italy	MIB 30	5	17.99	1.92
Japan	NIKKEI 225	20	15.59	0.81
UK	FTSE 100	10	14.62	3.62
US	S&P 500	25	15.68	1.66

The payoff of the option is

$$\max\left(0, \sum_{i=1}^{7} w_i \frac{S_{iT}}{S_{i0}} - 1\right),$$

where the w_i are weights that sum to 1. The correlation matrix of the 7 indices was erroneously printed in the paper and is taken from the erratum (Milevsky and Posner (1998b)). For easier reference we reproduce the matrix in Table 4. We set $r = 0.063$ and maturity $T = 1, 3, 5, 10$.

Table 4. Correlation matrix for the ILGIC indices

	Canada	France	Germany	Italy	Japan	UK	US
Canada	1.00	*	*	*	*	*	*
France	0.35	1.00	*	*	*	*	*
Germany	0.10	0.39	1.00	*	*	*	*
Italy	0.27	0.27	0.53	1.00	*	*	*
Japan	0.04	0.50	0.70	0.46	1.00	*	*
UK	0.17	−0.08	−0.23	−0.22	−0.29	1.00	*
US	0.71	0.15	0.09	0.32	0.13	−0.03	1.00

The left side of Table 5 shows Milevsky and Posner's results, and in the right we provide our MMC estimates, obtained using the same parameters and number of simulations ($N = 100000$).

Table 5. The left side of the table presents results taken from Milevsky and Posner (1998a), with standard deviations in brackets. The right side shows the MMC results

	Milevsky–Posner			MMC
T (years)	Plain MC	Rec. Gamma	Lognormal	
1	0.0587 (0.0002)	0.0589	0.0591	0.0590 (0.00007)
3	0.1331 (0.0004)	0.1328	0.1338	0.1335 (0.00009)
5	0.1945 (0.0005)	0.1942	0.1957	0.1952 (0.00009)
10	0.3104 (0.0007)	0.3103	0.3119	0.3113 (0.00005)

First of all, it is apparent from the SD's in brackets that MMC is much more accurate than plain MC methods for given N. In terms of computation times, this is equivalent to speed up calculations by 9 to 200 times. Secondly, we can confirm on a sound statistical basis the observation, made by Milevsky and Posner themselves, that the reciprocal gamma and lognormal approximations are downward and upward biased, respectively. The time needed to obtain one of the above MMC estimates is 3.3 seconds on an iBook 300 MHz. Given the rapidly growing computation power, we see little advantage in using a biased pricing formula when almost real-time estimates can be obtained by MMC.

3.3.2. Low dimensional portfolios We now present two simulation studies about portfolio options. We denote by $x \sim U[a, b]$ the statement that x is uniformly distributed on $[a, b]$. We first price, in the spirit of Broadie and Detemple (1996), a set of 1000 portfolio options with payoff

$$\max(0, q_1 S_{1T} + \ldots + q_n S_{nT} - k).$$

Assume that the correlation matrix of the log-returns of assets is of the form (5), $\rho \sim U[0, 1]$, the dimension $2 \leq n \leq 6$ is a uniformly sampled integer, the maturity is $T \sim U[0.25, 2.5]$, $S_{10} = \ldots = S_{n0} = 10$, $r \sim U[0, 0.1]$, no dividends are paid, each volatility is uniformly sampled in $[0, 0.5]$ and $q_i \sim U[0.5, 2], i = 1, \ldots, n$; provided $K = \sum_i^n q_i S_{i0}$, we set $k \sim KU[0.8, 1.2]$. We preliminarily evaluate the price of each option by 500000 antithetic simulations and discard the item whenever the price $\hat{C}_T \leq 0.5$, thus retaining $m = 970$ prices.

We then compute four estimates by $N = 16384$ simulations, namely, using antithetic MC, MMC and quasi-MC methods based on Sobol sequences (SMC). The last price is obtained using as control a European option on a single asset (hence the name 1MC), in order to evaluate the incremental contribution of MMC over this standard practice[1]. SMC denotes a deterministic integration scheme, which proved to be extremely effective in various financial pricing problems (see Joy et al. (1996)). For a given function of bounded variation, quasi-MC methods have errors going to zero as $O(\log^n N/N)$, which is close to $O(1/N)$ for low dimension n. For each of the three methods and for each of the 970 prices, we evaluate the percentage error, and summarize by taking the root mean squared error. For example, consider the SMC price: the relative root mean squared error is given by

$$\text{RMSE(SMC)} = \sqrt{\frac{1}{m}\sum_{i=1}^{m}\left(\frac{\hat{C}_T - \hat{C}_T(\text{SMC})}{\hat{C}_T}\right)^2},$$

where $\hat{C}_T$ is obtained by 500000 AVMC simulations and $\hat{C}_T(\text{SMC})$ denotes the estimate obtained by 16384 quasi-MC simulations. Likewise, we can define the RMSE for the other methods.

The use of RMSE instead of SD to measure errors appears to be appropriate when deterministic integration methods are used. In fact, the SD of quasi-MC methods is very often large in comparison to “truly” stochastic methods. Nevertheless, the estimated price is in many cases extremely accurate, and this explains the spreading use of RMSE (Dupire (1998)). Indeed, the use of SD is justified by *stochastic* large numbers theory that cannot be invoked to assess error properties of *deterministic* methods.

Table 6 shows the RMSE percentage errors for the four methods. On the one hand it is apparent that SMC prices are very close to the benchmark, giving average percentage error of 0.15%. On the other hand, MMC is always more accurate than AMC and 1MC methods. Apparently, MMC consistently provides worse results than SMC if percentage RMSE is considered. Even taking into account that SMC is about 45% slower[2] than MMC for given N, we find that quasi-MC methods appear to be very efficient. The only drawback of SMC might lie in the lack of precise error bounds or in the difficulties of performing statistical tests. See the conclusion of Morokoff (1998) for further insight on the advantages of the control variate approach over quasi-MC methods in some cases.

---

[1] In detail, we use as single control the European option on the asset such that $q_i\sigma_i$ is maximum. This is likely to select one of the most useful controls among the n assets and can be regarded as a heuristic selection of the set A in (4). We also tried to use a randomly chosen option or the one written on the asset corresponding to maximum q_i: in both cases the results are worse than those shown and this is true in the following Table 7 as well. All the values are obtained, as for MMC, with optimized coefficients.

[2] Our experience shows that this is mainly due to expensive accurate numerical inversion of cumulative normal (see Moro (1995)) instead of the widely used Box–Muller transformation that is deemed inappropriate for deterministic sequences.

Table 6. RMSE of estimated prices for up to 6-dimensional options

	RMSE ($\times 10^{-4}$)			
n	SMC	MMC	AMC	1MC
2	15	26	77	47
3	14	41	71	54
4	15	48	75	62
5	18	61	75	83
6	14	65	76	91
Total	15	50	75	69

Figure 1 shows a crossplot of the SD's of SMC against MMC prices. It is perhaps puzzling that the SD's of SMC prices are *always* bigger than those of MMC prices, especially after we note in Table 6 that very accurate estimates are provided by Sobol integration.

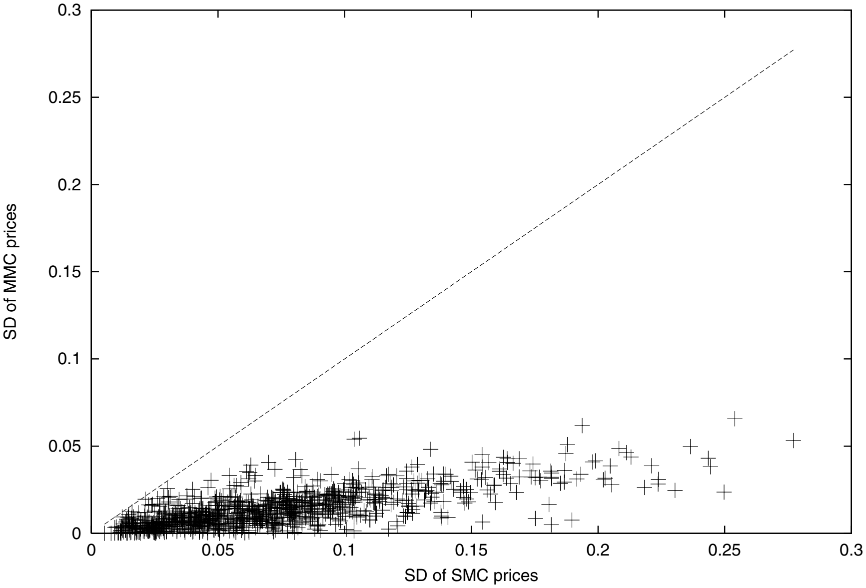

Fig. 1. Standard deviations of SMC and MMC prices, for up to 6-dimensional 970 portfolio options

As noted previously, the bare comparison of SD or RMSE might not be able to fully reveal savings in computing time. In order to measure the time efficiency of the methods, we assume that the errors all scales as $O(1/\sqrt{N})$ and evaluate the following speed-up factor for method W, using method Y

as a benchmark ($W, Y \in \{\text{MMC, AMC,1MC}\}$):

$$Su(W, Y) = \frac{\sigma_Y^2}{\sigma_W^2}\frac{t_Y}{t_W}, \tag{6}$$

where $\sigma_W, \sigma_Y, t_W, t_Y$ denote SD's and computation times of methods W and Y (and N is fixed). If $Su > 1$ then W is better than Y (in terms of standard deviation of the estimated price), provided that we let both methods run for the same time. Reciprocally, the ratio $1/Su$ is the time needed for W to evaluate one price with the same SD that Y achieves in 1 unit of time. The use of performance measures like Su or $1/Su$ is preferable as it is independent of CPU speed and so avoids the rapid obsolescence of CPU times that are often reported. Note that we cannot include SMC in the comparison as its error does not scale as required.

In Table 7 we give the quantity

$$S(W, Y) = \sum_{i=1}^{m} \frac{1}{Su(W, Y)_i},$$

where $Su(W, Y)_i$ is the speed-up factor obtained in pricing the i-th option, for different dimensions of portfolios.

Table 7. Times required by MMC, AMC, 1MC to estimate prices with the same standard deviation for up to 6-dimensional options. In squared brackets S is reported after the removal of 5 outlying data points

n	# samples	S(MMC,AMC)	S(AMC,AMC)	S(1MC,AMC)
2	180	18.37	180	232.20 [173.75]
3	227	41.70	227	237.50 [203.59]
4	188	43.70	188	606.57 [265.77]
5	166	48.47	166	276.25 [276.25]
6	209	71.28	209	451.09 [293.67]
Total	970	223.49	970	1803.59 [1213.04]

Table 7 shows, for example, that MMC can provide estimated prices with the same SD as AMC for 180 2-dimensional portfolio options in 18.37 units of time instead of 180 for AMC or 232.20 for 1MC. This is equivalent to saying that MMC is about 10 times faster than AMC in this case. It is apparent that performance with respect to AMC is decreasing with dimension, with MMC just 2.93 times faster for 6-dimensional options. Overall, AMC is 4.34 times slower than MMC and, if we set $W = \text{MMC}$ and $Y = \text{AMC}$, the factor Su is bigger than 1 in 953 out of 970 cases (98.2%). As far as 1MC is concerned, it is almost always considerably less efficient than AMC and the superior performance of MMC is striking. Note, however, that a tiny fraction of bad experiments badly affects the performance of 1MC: if we remove 5 outlying data points out of 970 then the results, shown in squared brackets in the table, are somewhat closer to AMC.

3.3.3. High dimensional portfolios The second application to portfolio options is on higher dimensional baskets, that are composed of $10 \leq n \leq 40$ assets. We compare AMC and MMC prices, using the speed-up factor (6).

All the parameters of the priced basket options are selected randomly, as previously done for the low dimensional options, the only difference being in greater n. We price 1000 portfolio options using 1000000 AMC simulations, discard 8 cases for which the computed prices was smaller than 0.5 and are then left with $m = 992$ valid samples. Figure 2 depicts the speed-up factors Su(MMC,AMC) against the value of the basket.

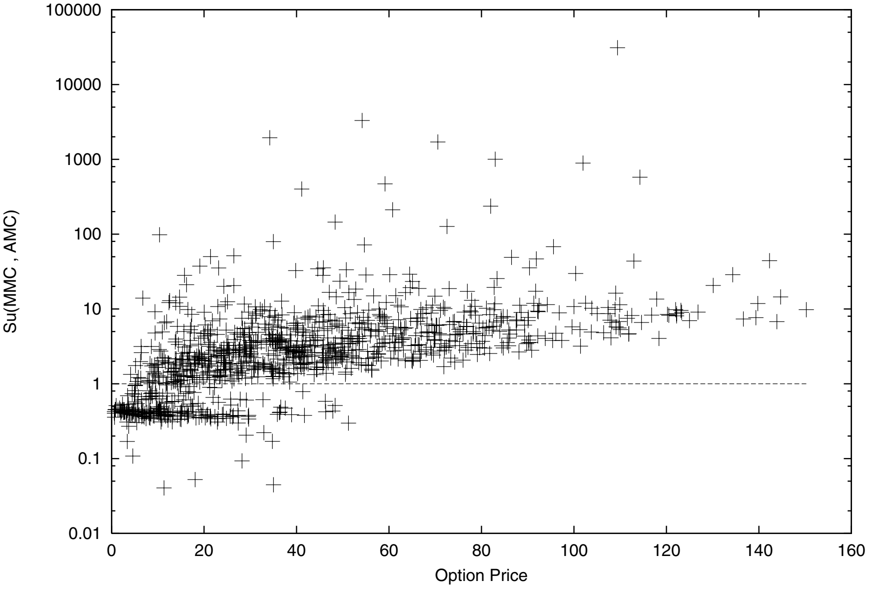

Fig. 2. Su factors of MMC as a function of the value of the portfolio, for up to 40-dimensional options. The constant $Su = 1$ is relative to AMC taken as benchmark. Points below the horizontal line denote options for which AMC is preferable to MMC

Figure 2 can be interpreted in at least two ways. On the one hand, counting the points above the benchmark AMC line $Su = 1$, we get the percentage of cases that are more efficiently priced by MMC; exact calculations show that they are 792 out of 992 (about 80%). On the other hand, we see that the efficiency of the method appears to increase with the value of the option. As there is a standard way to improve performance for out-of-the-money options, namely, importance sampling (see, for example, Fishman (1996) or Glasserman et al. (1999) for a recent application), this technique could be used together with MMC to get the most out of both methods. Examination of other graphs like Fig. 2, not reported for brevity, shows that the performance is quite insensitive to time to expiration T, to risk-free rate and correlation among assets.

Table 8 shows $S(\text{MMC,AMC})$ and $S(\text{AMC,AMC})$. It is seen that a modest 19.5% speed improvement against AMC is achieved. Hence, recalling the quite different picture in Section 3.3.2, we can deduce that the efficiency of MMC is decreasing with dimension n, a very common feature of other methods as well. In the very few cases where proper estimation of the coefficient vector **b** is difficult due to ill-conditioned regression, we set $b_1 = b_2 = \ldots = b_n = 1$. The computation times required to evaluate S are inclusive of all these numerical subtleties and therefore still offer a fair comparison among methods. We are not providing results for 1MC as it is almost 5 times less efficient than AMC.

Table 8. Times required by MMC and AMC to estimate prices with the same standard deviation for up to 40-dimensional options

n	# samples	$S(\text{MMC,AMC})$	$S(\text{AMC,AMC})$
10–20	349	245.99	349
21–30	314	255.21	320
31–40	337	296.88	323
Total	992	798.08	992

3.4. *"Multiple" options*

In this section we compare the SD of plain MC, AMC and MMC on "multiple" options (also called outperformance options) having payoffs $\max(0, S_{1T} - k_1, \ldots, S_{nT} - k_n)$, where the dimension n varies from 2 to 10. All the parameters are randomly selected as for portfolio options. The strikes k_i are such that $k_i \sim S_{i0}U[0.8, 1.2]$, $i = 1, \ldots, n$. Again we price 1000 "multiple" options and are left with $m = 991$ valid samples after elimination of the options whose AMC price $\hat{C}_T < 0.5$ (using 1000000 simulations). We compare MC, AMC, MMC and MMC(i_1, i_2) as described below. In view of the final considerations of Section 2, we select the set A of control variates by the following rule of thumb: we use just two control variates, the first relative to the asset whose strike is minimum and the second relative to the asset with maximum volatility. Hence we selected just

$M_T(i), i \in A = \{i_1, i_2\}$ where

$$i_1 = \arg\min_i k_i, \qquad i_2 = \arg\max_i \sigma_i.$$

Note that this is far from being an optimized search for the optimal set of control variates that we mentioned above, but is only intended to show that accurate results can be obtained with a few carefully chosen $M_T(i)$'s.

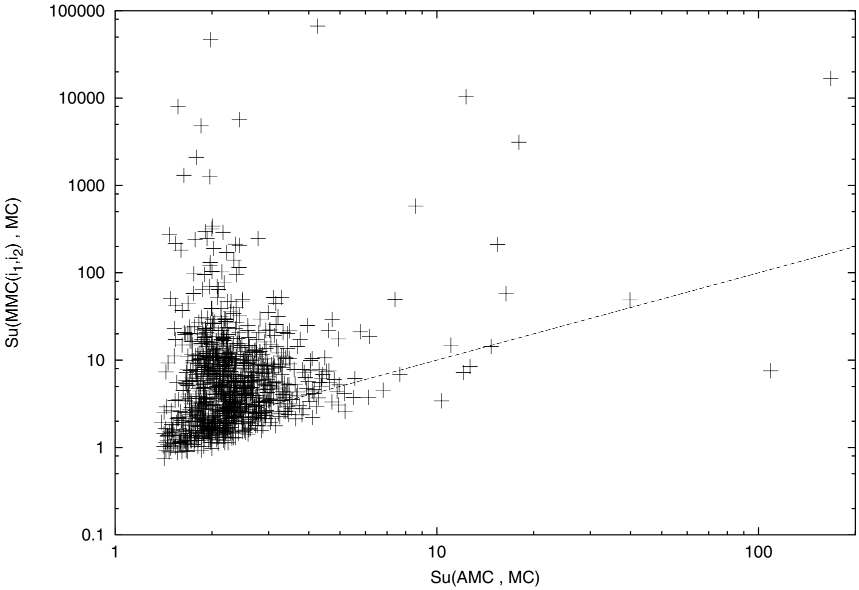

Fig. 3. Su factors of AMC against MMC(i_1, i_2) for multiple options

Figure 3 shows the crossplot of $Su(\text{MMC}(i_1, i_2), \text{MC})$ against Su(AMC,MC). The graph shows that the use of MMC(i_1, i_2) is extremely useful on many occasions, while slightly slower than AMC in the remaining cases. In detail, MMC(i_1, i_2) has greater Su factor than AMC in 742 out of 991 cases (74.9%). Table 9 shows the usual indicators $S = \sum_{i=1}^{m} 1/Su_i$ for MC, AMC, MMC and MMC(i_1, i_2), grouped according to dimension. The overall fastest method is MMC(i_1, i_2) which prices the 991 options in 288.37 time units, 3.44 times faster than MC and 35% more effective than AMC.

The MMC results are very good for $n \leq 7$, but using all the controls underperforms both MMC(i_1, i_2) and AMC for $8 \leq n \leq 10$. Figure 4 depicts S against n for plain MC, MMC, MMC(i_1, i_2) and AMC and shows that performance is in general decreasing with dimension. We think that these results mainly show the potential of an optimal search for the optimal set A, that hopefully will be implemented in future research.

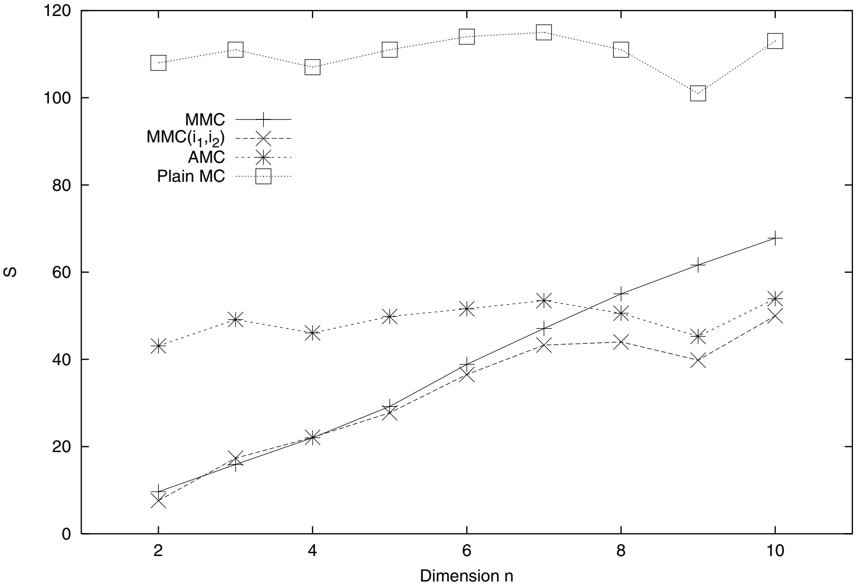

Fig. 4. S against dimension n for plain MC, MMC, AMC, MMC(i_1, i_2) applied to multiple options

Table 9. Times required to estimate prices with the same standard deviation as MC for up to 10-dimensional multiple options

n	# samples	S(MMC,MC)	S(MMC(i_1, i_2),MC)	S(AMC,MC)
2–4	326	47.56	47.05	138.27
5–7	340	115.12	107.53	154.91
8–10	325	184.45	133.80	149.77
Total	991	347.13	288.37	442.95

4. Conclusion

We presented a variance reducing control variate procedure that can price a host of European multivariate options, which are important in various areas of asset management and corporate decision making. Our method takes advantage of the existence of pricing formulae for univariate options, replacing all but one random asset in the payoff function by their expectations. We called the resulting technique Mean Monte Carlo (MMC).

We tested the methodology on a variety of multi-asset options, proving its broad applicability and accuracy against closed evaluation formulae or other numerical pricing methods, both of Monte Carlo type and based on different ideas such as density approximation. We also performed three large sample simulation studies on various kind of options whose dimension ranges from 2 to 40, comparing our results to plain, antithetic and quasi-MC methods.

In detail, after pricing an exchange option we showed that the Milevsky and Posner approximation is slightly biased for portfolio options and that an accurate price of an ILGIC can be provided in a few seconds. Some extreme examples taken from Barraquand's papers were used to compare MMC to the quadratic resampling method, resulting in more accurate estimates in 15 out of 22 cases. Overall, MMC is 4.34 faster than AMC for low dimension portfolio options, and of comparable speed if baskets of high dimension (less than 40) are considered. Though MMC can be used to improve quasi-MC as well, we gave empirical evidence that deterministic MC is more accurate than stochastic MC methods if the dimension is low. In pricing "multiple" options we proposed a rule of thumb to select two control variates that achieved an average 35% increment in speed against AMC. This suggests that further improvement in variance reduction could be obtained by optimal choice of control variates.

In general, MMC appears to perform best in high volatility environments and in pricing in-the-money options. This suggests that combined use of importance sampling and MMC should be investigated. MMC results appear to be very good for low to medium dimension but degrade somewhat on high dimensional options, which are priced with similar effectiveness by AMC.

References

Barraquand, J. (1993): Numerical valuation of high dimensional European securities. Research Report **26**. Digital Equipment Corporation Paris Research Laboratory, Paris

Barraquand, J. (1995): Numerical valuation of high dimensional multivariate European securities. *Management Science* **41**, 1882–1891

Boyle, P., Broadie, M., Glasserman, P. (1997): Monte Carlo methods for security pricing. *Journal of Economic Dynamics and Control* **21**, 1267–1321

Boyle, P.P. (1977): Options: a Monte Carlo approach. *Journal of Financial Economics* **4**, 323–338

Broadie, M., Detemple, J. (1997): Recent advances in numerical methods for pricing derivative securities. In: Rogers, L.C.G., Talay, D. (eds.): Numerical methods in finance. Cambridge University Press, Cambridge, pp. 43–66

Cox, J.C., Ross, S.A. (1976): The valuation of options for alternative stochastic processes. *Journal of Financial Economics* **3**, 145–166

Cox, J. C., Ross, S. A., Rubinstein, M. (1979): Option pricing: a simplified approach. *Journal of Financial Economics* **7**, 229–263

Dupire, B. (ed.) (1998): Monte Carlo: methodologies and applications for pricing and risk management. Risk Books, London

Fishman, G.S. (1996). Monte Carlo: concepts, algorithms, and applications. Springer, New York

Glasserman, P., Heidelberger, P., Shahabuddin, P. (1999): Importance sampling in the Heath-Jarrow-Morton framework. *The Journal of Derivatives*, **7**, 32–50

Hammersley, J. M., Handscomb, D.C. (1964): Monte Carlo methods. Methuen, London

Joy, C., Boyle, P., Tan, K. (1996): Quasi-Monte Carlo methods in numerical finance. *Management Science* **42**, 926–938

Kemna, A.G.Z., Vorst, A.C.F. (1990): A pricing method for options based on average asset values. *Journal of Banking and Finance* **14**, 113–129

Krommer, A., Ueberhuber, C. (1998): Computational integration. SIAM, Philadelphia

Margrabe, W. (1978): The value of an option to exchange one asset for another. *The Journal of Finance* **33**, 177–186

Milevsky, M., Posner, S. (1998a): A closed-form approximation for valuing basket options. *The Journal of Derivatives* **5**, 54–62

Milevsky, M., Posner, S. (1998b): Erratum. A closed-form approximation for valuing basket options. *The Journal of Derivatives* **5**, 83

Moro, B. (1995): The Full Monte. *Risk* **8**, 57–58

Morokoff, W. (1998): Generating quasi-random paths for stochastic processes. *SIAM Review* **4**, 765–788

Pellizzari, P. (2001): Efficient Monte Carlo pricing of portfolio options. *Rendiconti per gli Studi Economici*, Quantitativi 2001, 108–123, http://www.dma.unive.it/∼paolop/papers/rendidef.pdf

Ross, S. (1997): Simulation. 2nd edition. Academic Press, London

Rubinstein, R. (1981): Simulation and the Monte Carlo method. Wiley, New York

Schmeiser, B., Taaffe, M., Wang, J. (2001): Biased control-variate estimation. IIE Transactions **33**, 219–228

Schmeiser, B., Taaffe, M., Wang, J. (2002): Control-variate estimation using estimated control means. IIE Transactions. To appear

Shaked, M., Shantikumar, G. (1994): Stochastic orders and their applications. Academic Press, Boston

Stulz, R. (1982): Options on the minimum or the maximum of two risky assets. *Journal of Financial Economics* **10**, 161–185

Wilmott, P., Dewynne, J., Howison, S.D. (1995): The mathematics of financial derivatives. A student introduction. Cambridge University Press, Cambridge