

REDUCING THE VARIANCE OF LIKELIHOOD RATIO GREEKS IN MONTE CARLO

Luca Capriotti

Global Modelling and Analytics Group
Credit Suisse Group, Investment Banking Division
Eleven Madison Avenue, New York, NY 10010, U.S.A.

ABSTRACT

We investigate the use of Antithetic Variables, Control Variates and Importance Sampling to reduce the statistical errors of option sensitivities calculated with the Likelihood Ratio Method in Monte Carlo. We show how Antithetic Variables solve the well-known problem of the divergence of the variance of Delta for short maturities and small volatilities. With numerical examples within a Gaussian Copula framework, we show how simple Control Variates and Importance Sampling strategies provide computational savings up to several orders of magnitude.

1 INTRODUCTION

Monte Carlo (MC) simulations are one the main tools employed in the Financial Services industry for pricing and hedging derivatives securities. In fact, as a result of the ever increasing level of sophistication of the financial markets, a considerable fraction of the pricing models employed by investment firms is too complex to be treated by analytic or deterministic numerical methods. For these models, MC simulation is the only computationally feasible pricing method.

The main drawback of MC methods is that they are generally computationally expensive. These efficiency issues become even more dramatic when MC simulations are used for the calculation of price sensitivities, i.e., the derivatives of the option price with respect to the parameters of the underlying model, also known as *Greeks*. In fact, the standard method for the calculation of a price sensitivity, say with respect to a parameter θ_k , is based on a finite difference approximation of the derivatives $\partial V(\theta)/\partial \theta_k$. This method, also known as ‘bump and reval’, involves repeating the MC simulation, and evaluating the finite difference estimate

$$\frac{\partial V(\theta)}{\partial \theta_k} \simeq \frac{V(\theta_1, \dots, \theta_k + \Delta\theta, \dots, \theta_n) - V(\theta)}{\Delta\theta}, \quad (1)$$

for a small increment $\Delta\theta$. The main virtue of this method is that it is straightforward to understand, and it requires minimal implementation effort. The drawback is that additional MC simulations are required for each sensitivity, and that the finite differences (1) may be affected by large statistical errors, especially for payout with discontinuities (Glasserman 2004). As a result, hedging derivative securities with MC simulations can be extremely time consuming.

Alternative methods for the calculation of price sensitivities have been proposed in the literature (for a review see Glasserman 2004). Here we concentrate on the so-called Likelihood Ratio Method (LRM). The principal advantage of this technique when compared to ‘bump and reval’ is that it allows to calculate all the sensitivities simultaneously in a single MC simulation, and a single set of payout evaluations. In addition, the variance properties of LRM estimators are not affected as much by discontinuities in the payoff. As a result, for digital and barrier options, LRM may provide a better convergence than bumping (Glasserman 2004). The main drawback is that the statistical uncertainties of LRM estimators are nonetheless difficult to predict, and can be sometimes large. What is worse, in some cases such uncertainties are even known to diverge thus making the MC simulation very time consuming if not hopeless in practice.

In order to address this difficulty, in this paper we investigate three Variance Reduction techniques – Antithetic Variables, Control Variates, and Importance Sampling – that can dramatically improve the MC convergence of LRM estimators. In the next Section, we begin by reviewing the rationale of LRM, specializing our discussion to a Gaussian Copula framework very common in the financial practice. The use of Antithetic Variables is discussed in Section 3. In particular, we will show how this simple technique solves the well-known problem of the divergence of the variance of LRM Deltas for short maturity and low volatility. Then, in Section 4 we illustrate how Control Variates and Importance Sampling can drastically suppress the statistical uncertainties of the LRM estimators thus reducing the computational cost for the Greeks by orders of magnitude.

2 LIKELIHOOD RATIO METHOD

The arbitrage-free price of a derivative security can be expressed in general as the expectation value of the discounted cash flows, $G(x)$, over a risk-neutral probability density (Harrison and Kreps 1979), $P_\theta(x)$,

$$V(\theta) = \mathbb{E}_P[G(x)] = \int dx G(x) P_\theta(x) , \quad (2)$$

where $x = (x_1, \dots, x_N)$ is a N -dimensional vector representing the underlying random factors upon which the claim is contingent. Here the vector $\theta = (\theta_1, \dots, \theta_n)$ represents a set of parameters whose value is generally determined by calibrating the chosen model or, equivalently, the density $P_\theta(x)$, on the prices of securities liquidly traded in the market.

Whenever the dimension N of the state variable x is large (say $N \gtrsim 4$) MC methods are the only feasible route for estimating expectation values of the form (2). In their simplest incarnation, these consist in averaging the payout function $G(x)$ over N_p independent random realizations of the vector x , say $x[m]$, generated according to the probability density $P_\theta(x)$,

$$V(\theta) \simeq \bar{V} = \frac{1}{N_p} \sum_{m=1}^{N_p} G(x[m]) . \quad (3)$$

In fact, the central limit theorem (Kallenberg 1997) ensures that, for big enough samples, the values of the estimator \bar{V} are normally distributed around the true value, and converge for $N_p \rightarrow \infty$ towards V namely

$$V \simeq \frac{1}{N_p} \sum_{m=1}^{N_p} G(x[m]) \pm \frac{\kappa}{\sqrt{N_p}} , \quad (4)$$

where $\kappa^2 = E_P[G(x)^2] - E_P[G(x)]^2$ is the variance of the MC estimator. Here, following a common terminology, we refer to $\kappa/\sqrt{N_p}$ as *statistical uncertainty* or *statistical error*. Although Eq. (4) ensures the convergence of the MC average to the expectation value (2) provided that κ is finite, the square root law in (4) can make the calculation of accurate estimates time consuming.

This is particularly true, for the MC calculation of the Greeks. In fact, the variance and bias properties of finite difference estimators of the form (1) can be in some cases rather poor (Glasserman 2004). This is because, while the bias of the finite difference (1) can be made in general arbitrarily small by reducing the value of $\Delta\theta$, its statistical error can in some common cases diverge for $\Delta\theta \rightarrow 0$. When this happens, choosing a value of $\Delta\theta$ small enough to reduce the bias to an acceptable level may require a large computational




Figure 1: Delta of a European Call option (21) for $S_0 = 100$, $K = 25$, and $r = 0.05$ as a function of the time to maturity T for $\sigma = 0.1$ (a), and as a function of the volatility σ for $T = 1/12$ (b): crude MC (triangles and continuous line), MC with Antithetic Variables (circles and dashed line).

cost in order to obtain statistically accurate results. For a convergence analysis see e.g., (Milstein and Tretyakov 2005).

Several methods have been recently proposed in the literature in order to speed up the calculation of option sensitivities (Glasserman 2004). Here we will concentrate on the so-called Likelihood Ratio Method (LRM). Under mild regularity conditions on the probability density $P_\theta(x)$, the sensitivity of the option price (2) with respect to any parameter θ_k can be obtained as

$$\bar{\theta}_k = \frac{\partial V(\theta)}{\partial \theta_k} = \mathbb{E}_P[G(x)\Omega_k(x)] , \quad (5)$$

i.e., by calculating the expectation value of the original payout function multiplied by the so-called *Likelihood Ratio* weight

$$\Omega_k(x) = \frac{\partial \log P_\theta(x)}{\partial \theta_k} , \quad (6)$$

giving as MC estimator:

$$\bar{\theta}_k \simeq \frac{1}{N_p} \sum_{m=1}^{N_p} G(x[m])\Omega_k(x[m]) . \quad (7)$$

In the following, we will specialize our discussion to the case where the probability distribution $P_\theta(x)$ is a N -dimensional Gaussian Copula. The latter is defined by a correlation matrix Σ , and a set of N marginal cumulative

distributions $M_i(x_i)$, so that the joint distribution reads

$$F(x) = \prod_{i=1}^N \int_{-\infty}^{x_i} dy_i P(y_1, \dots, y_N) = \Phi_N(\Phi^{-1}(M_1(x_1)), \dots, \Phi^{-1}(M_N(x_N)); \Sigma). \quad (8)$$

Here $\Phi_N(x_1, \dots, x_N; \Sigma)$ is the cumulative distribution of a N -dimensional Gaussian random variable with zero mean and correlation matrix Σ , and $\Phi(\cdot)$ is the standard normal cumulative distribution.

If we restrict to the case in which the correlation matrix Σ does *not* depend on the parameters θ , the LRM weight (6) for the Gaussian Copula distribution (8) can be expressed as

$$\Omega_\theta(x) = \sum_{i=1}^N \partial_\theta \log m_i(x_i) - Z(x)^T (\Sigma^{-1} - I) \partial_\theta Z(x), \quad (9)$$

where I is the N -dimensional identity matrix, $m_i(x_i) = dM(x_i)/dx_i$ is the i -th marginal density function, and the components of the vector $Z(x)$ are

$$Z_i = \Phi^{-1}(M_i(x_i)), \quad (10)$$

and those of $\partial_\theta Z(x)$ read

$$\partial_\theta Z_i = \frac{\partial_\theta M_i(x_i)}{\phi(\Phi^{-1}(M_i(x_i)))}, \quad (11)$$

with $\phi(\cdot)$ the standard normal density function. A derivation of Eq. (9) is given in the Appendix.

Gaussian Copula models of the form (8) are widely used in Financial Engineering. Indeed, the above formulation can be used to evaluate structured European options written on several assets, e.g., equity, commodity, rates or foreign exchange pairs (Hull 2002). In this case, the marginal distribution of each asset is typically implied from liquidly traded Vanilla options, and the co-dependence between the factors is modeled by means of the Gaussian Copula. Structured credit pricing, e.g., for CDO and CDO² (Schonbucher 2003), can be also performed within a similar framework. In general, whenever the marginal distributions above are not known in closed form, e.g., they are calculated numerically by means of a calibration procedure, the derivatives in Eq. (11) can be easily computed by means of finite differences. This does not generally introduce accuracy or stability problems provided the calibration algorithms employed are numerically stable.

It is easy to see that the LRM weights above give the expected result in the case a multi-asset lognormal model of the form

$$S_i = S_i^0 \exp[(r - \sigma_i^2/2)T + \sigma_i \sqrt{T} Z_i], \quad (12)$$

where S_i^0 and σ_i are the spot price and volatility of the i -th asset, T is the maturity of the option, and Z_i are standard normal variables with correlation $\Sigma_{ij} = \mathbb{E}[Z_i Z_j]$. In fact, one clearly has $x_i = Z_i$ with

$$Z_i = \frac{\log S_i/S_i^0 - (r - \sigma_i^2/2)T}{\sigma_i \sqrt{T}} \quad (13)$$

so that the LRM weights for the i -th Delta and Vega (Hull 2002) read respectively

$$\Omega_\Delta^i(Z) = \frac{\partial}{\partial S_i^0} \log P_\theta(x) = \frac{[\Sigma^{-1} Z]_i}{\sigma_i \sqrt{T} S_i^0}, \quad (14)$$

and

$$\Omega_V^i(Z) = \frac{\partial}{\partial \sigma_i} \log P_\theta(x) = \left(\frac{Z_i}{\sigma_i} - \sqrt{T}\right) [\Sigma^{-1} Z]_i - \frac{1}{\sigma_i}. \quad (15)$$

It is straightforward to verify that these equations are in agreement with the general expression for the LRM weight in a Gaussian model given in (Glasserman 2004)

$$\Omega_\theta(Z) = -\frac{1}{2} \text{Tr}[\hat{\Sigma}^{-1} \partial_\theta \hat{\Sigma}] + \frac{1}{2} X \hat{\Sigma}^{-1} (\partial_\theta \hat{\Sigma}) \hat{\Sigma}^{-1} X + X \hat{\Sigma}^{-1} \partial_\theta m, \quad (16)$$

with $m_i = \log S_i^0 + (r - \sigma^2/2)T$, $X_i = \sigma_i \sqrt{T} Z_i$, and $\hat{\Sigma}_{ij} = \sigma_i \sigma_j \Sigma_{ij}$.

In the special case of a single asset, the weights above simplify to the well-known expressions

$$\Omega_\Delta(Z) = \frac{Z}{\sigma \sqrt{T} S^0}, \quad (17)$$

and

$$\Omega_V(Z) = \frac{Z^2 - 1}{\sigma} - Z \sqrt{T}. \quad (18)$$

3 SOLVING THE PROBLEM OF DELTA'S DIVERGING VARIANCE WITH ANTITHETIC VARIABLES

For a given number of MC iterations the calculation of the Greeks by means of LRM is generally fast when compared to bumping. However, the speed of convergence of the LRM estimators is difficult to predict *a priori* for a given problem, it is payout and parameters dependent, and can be in some instances particularly slow (Glasserman 2004). In fact, since the LRM weight has in general zero mean as a result of the identity

$$\partial_\theta \int dx P_\theta(x) = 0, \quad (19)$$

the LRM estimators for the Greeks have no definite sign. This can give rise to poor variance properties whenever the configurations with opposite sign have similar weight in the MC average (7) so that the final outcome is the result of the cancellation of two comparable and not necessarily highly correlated quantities.

Table 1: Delta of the Basket Call option (27) for $T = 0.5$ and $r = 0.05$. The volatilities of the assets are all equal to $\sigma = 0.3$, and their Forwards range between 51.3 and 55.9. The uncertainties are reported in parenthesis.

K	AV	AV+CV	AV+LSIS
30	7.5(5)	7(1)10 ⁴	1200(100)
40	3.5(3)	1200(100)	200(10)
50	2.2(2)	410(3)	100(10)
60	2.1(1)	11(1)	120(10)
70	2.2(2)	3.1(3)	200(20)
80	1.9(2)	2.1(3)	610(70)

In particular, a common problem generally reported in the literature (Glasserman 2004, Jaeckel 2002) is the divergence of the variance of the LRM weight for Delta Eqs. (14) and (17) in the limit of small volatility and short maturity. Indeed, from Eq. (14)

$$\text{Var}[\Omega_{\Delta}^i] \propto \frac{1}{\sigma_i^2 T} \rightarrow \infty \quad (20)$$

for $\sigma_i \sqrt{T} \rightarrow 0$. This is illustrated in Figure 1 for a ‘deep in the money’ Call option (i.e., with very low strike compared to the expectation value of the underlying asset, or Forward, see Hull 2002) on a single lognormal asset with undiscounted payout

$$G(S) = (S - K)^+ \quad (21)$$

where K is the strike price. The LRM estimator becomes extremely noisy for $T \rightarrow 0$ and for small volatility, to the point of not providing any useful information for maturities shorter than a few weeks for any practical number of MC iterations.

The divergence of the LRM weight for Delta is due to the break down of the absolute continuity property of the probability density function, which is required to take the derivative inside the expectation in Eq. (5) (Glasserman 2004). This generally affects the LRM estimators (9) for Delta also for non-lognormal models.

Although, to the best of our knowledge, it was not previously noted in the literature, this problem can be easily overcome by using Antithetic Variables. Indeed, since the LRM weight (14) is *odd* in each of the Gaussian random variable Z_i , the Antithetic estimator (Glasserman 2004) for




Figure 2: Vega of the Basket Call option (27) as a function of the strike price. Crude MC: triangles and dashed line. MC with Variance Reduction techniques: circles and continuous line. Empty circles: Antithetic Variables with Control Variates. Full Circles: Antithetic Variables with LSIS.

the weight reads

$$\Omega_{\Delta}^i|_{ant} = \frac{\Omega_{\Delta}^i(Z) + \Omega_{\Delta}^i(-Z)}{2} \equiv 0 \quad (22)$$

so that $\mathbb{E}_P[\Omega_{\Delta}^i|_{ant}] = 0$, with *zero variance*. What is more, it is also possible show that the variance of the product of the payout and the weight in (7) is generally bounded as $\sigma \sqrt{T} \rightarrow 0$. This can be realized by means of the following simple heuristic argument. Consider for simplicity a single asset payout under the lognormal model (12). This can be approximated for small maturities and volatilities as

$$P(Z) \simeq c_0 + c_1 \sigma \sqrt{T} Z + O(\sigma^2 T) \quad (23)$$

with c_0 and c_1 constants. As a result, the LRM Antithetic estimator for Delta (7) reads in this limit

$$\Omega_{\Delta}|_{ant} \simeq c_0 \frac{Z - Z}{2\sigma\sqrt{T}S^0} + c_1 \frac{Z^2}{S^0} + O(\sigma\sqrt{T}) \quad (24)$$

whose variance is clearly bounded for $\sigma \sqrt{T} \rightarrow 0$ as a result of the cancellation of the leading term.

Figure 1 illustrates the efficacy of the method: the Antithetic LRM estimator provides a stable Delta and a practically constant statistical uncertainty for all maturities and volatilities. The variance reduction with respect to the crude LRM estimators is around 10 for 1 year, and around 500 for 1 week maturity.

While the results presented here are for a lognormal model, we have found comparable variance reductions also for market-implied marginal distributions $M_i(x)$ for a variety of equity, foreign exchange, interest rate and commodity underlying assets. Indeed, the antisymmetry of the LRM

weight (9) for Delta is generally satisfied, at least approximately, also for skewed distributions. It is also worth noting that although here we have limited our discussions to European-style payouts in a Gaussian Copula framework, it is easy to realize that Antithetic Variables generally solve the problem of the divergence of the variance of Delta estimators also for path-dependent options (similar ideas have been also used by Mike Giles in the implementation of the so-called Vibrato Monte Carlo technique (Giles 2007)).

4 CONTROL VARIATES AND IMPORTANCE SAMPLING

Although the variance of LRM estimators generally remains finite when using Antithetic Variables, their convergence can still be poor in some cases, especially for Vega. For this reason, an efficient implementation of LRM generally requires the use of other Variance Reduction strategies. Here we consider Control Variates and Importance Sampling (Glasserman 2004).

The Control Variates method aims to reduce the statistical uncertainty of a MC average by exploiting the correlation of its statistical samples with those of some quantity (the control) whose expectation value is known *a priori* (Glasserman 2004). This technique may result in spectacular variance reductions but requires some closely related estimator with a known integral over the sampled probability distribution.

For this discussion, we will restrict ourselves to controls that in our context are generally readily available, and we will consider the LRM weights and the derivatives of the first moment of the marginal distributions in Eq. (8). In fact, as previously mentioned, the LRM weight (6) has always zero expectation value. It is also reasonable to expect the weight to be somewhat correlated with the corresponding Greek estimator in (7). On the other hand, the first moment of each marginal distribution is usually known as it contains the information on the Forward of the underlying asset, which is the first thing to be usually matched with the available market data. Its derivatives with respect to the main model parameters are also generally known. For instance, in a typical financial context, including models with skew (Hull 2002), one has

$$\frac{\partial \mathbb{E}[S_i]}{\partial S_i^0} = P(0, T)^{-1} , \quad (25)$$

where $P(0, T)$ is today's ($t = 0$) price of a zero coupon bond maturing at time T , and

$$\frac{\partial \mathbb{E}[S_i]}{\partial \sigma_i} = 0 . \quad (26)$$

In the following, when calculating the Greeks of the i -th asset, we will use the corresponding LRM weight, and Risk of the Forward. Note, however, that when calculating the LRM Delta with Antithetic Variables, we can only use the Delta of the Forward as a control, as the Antithetic estimator for the LRM weight (22) is identically zero.

Table 2: Same of Table 1 for Vega.

K	AV	AV+CV	AV+LSIS
30	1.3(1)	4(1)10 ⁴	11(1)
40	1.7(1)	150(10)	7.0(6)
50	2.1(1)	22(2)	8.0(8)
60	2.1(2)	5.5(6)	80(10)
70	1.9(2)	5.2(5)	340(40)
80	1.7(3)	1.7(3)	1100(100)

Importance Sampling techniques, on the other hand, do not rely on the knowledge of any closely correlated estimator but aim to reduce the variance by sampling more effectively the domain of integration in Eq. (2). Here we will use a recently introduced Importance Sampling strategy based on a Least-squares optimization, namely the Least-Squares Importance Sampling (LSIS) (Capriotti 2008, Capriotti 2007). In particular we will use as trial densities mean-shifted single mode and bi-mode multivariate Gaussian distributions (see references above).

In Table 1, we compare the results obtained by using Antithetic Variables only, and in combination with Control Variates and Importance Sampling for the Delta of one of $N = 10$ assets of a Basket Call option with undiscounted payout

$$G(S) = \left(\frac{1}{N} \sum_{i=1}^N S_i - K \right)^+ , \quad (27)$$

in a lognormal model of the form (12). Here, as an indicator of the efficiency gains introduced by the different methods, we have defined the variance (efficiency) ratio as

$$ER = \left(\frac{\sigma[\text{CrudeMC}]}{\sigma[\text{VR}]} \right)^2 \quad (28)$$

where the numerator and denominator are respectively the statistical errors (for the same number of MC paths) of the crude estimator and of the one obtained with different Variance Reduction techniques.

As shown in Table 1, for the considered maturity, the efficacy of Antithetic Variables decreases moving away from the deep in the money region, and is generally limited to around a factor of two for larger strikes. Using the Delta of the corresponding Forward (25) as Control Variate provides about an order of magnitude variance reduction around the 'at the money' point (i.e., for strikes around the Forward

value). Decreasing further the strike results in a spectacular suppression of the statistical errors of the Control Variate estimator. This is expected due to the very high correlation between the option payoff in this regime and a simple Forward contract. However, as the strike moves instead in the ‘out of the money’ region (high strikes compared to the Forward) such correlation rapidly decreases, and Control Variates become practically ineffective.

On the other hand, a simple Importance Sampling strategy based on a single mean-shifted Gaussian trial density and implemented by means of the LSIS approach proves to be very effective at all level of moneyness. In fact, although Importance Sampling gives smaller efficiency gains than Control Variates for deep in the money options, it provides at least two orders of magnitude speed up across all strikes, including the out of the money regions where Control Variates lose their efficacy. Indeed, as it can be generally expected when using mean-shifted trial densities (Capriotti 2008), Importance Sampling is particularly effective moving away from the money in either direction.

LRM estimators for Vega are generally noisier than the ones for Delta. This can be understood from the form of the weight in Eq. (15) as it involves the second moments of the Random increments which have generally a larger variance. This is illustrated in Fig. 2 where we plot one of the Vegas for the Basket option above as a function of the strike: for low strikes, as Vega becomes smaller, the crude MC estimate becomes extremely noisy. As mentioned, this is due to fact that the small value of Vega in this regime is the result of the cancellation of two poorly correlated stochastic quantities representing the averages of the LRM estimator over the configurations in which it has a definite (positive and negative) sign.

As expected, due to the presence in the LRM weight of even terms in the random increments, the efficacy of Antithetic Variables is very limited for Vega (see Table 2). In contrast, using both the LRM weight and the Vega of the Forward as Control Variates provides sizable variance reductions both for in the money and at the money options, with the expected remarkable efficiency gains for small strikes. On the other hand, Importance Sampling – although producing smaller variance reductions for in the money options – becomes particularly effective for larger strikes.

In this sense, Importance Sampling and Control Variates appear to be somewhat complementary for this problem, and Figure 2 displays the LRM Vega as obtained by choosing the most effective of the two techniques for each strike. This results in orders of magnitude savings in computer time with respect to the crude MC calculation.

ACKNOWLEDGMENTS

It is a pleasure to acknowledge useful discussions with Mike Giles, Paul Glasserman, David Shorthouse, Jacky

Lee, Anton Merlushkin, Mark Stedman, and Sanjay Chawla. The opinions and views expressed in this paper are uniquely those of the author, and do not necessarily represent those of Credit Suisse Group.

A APPENDIX: DERIVATION OF EQ. (9)

First, by differentiating the joint cumulative distribution (8) one obtains the corresponding probability density function

$$P(x) = \phi_N(\Phi^{-1}(M_1(x_1)), \dots, \Phi^{-1}(M_N(x_N)); \Sigma) \times \prod_{i=1}^N \frac{m_i(x_i)}{\phi(\Phi^{-1}(M_i(x_i)))}, \quad (29)$$

where $m_i(x_i) = dM(x_i)/dx_i$ is the i -th marginal density function, and $\phi_N(x_1, \dots, x_N; \Sigma)$ is the multivariate Gaussian density with correlation Σ . Then, taking the logarithm of Eq. (29) gives

$$\log P(x) = \sum_{i=1}^N \left(\log m_i(x_i) - \log \phi(Z_i) \right) - \frac{1}{2} Z(x)^T \Sigma Z(x) - \frac{N}{2} \log 2\pi - \frac{1}{2} \log (\det \Sigma), \quad (30)$$

where we have used the explicit form of $\phi_N(Z_1, \dots, Z_N; \Sigma)$, and the definition in Eq. (10). Hence, the derivative with respect to θ of the latter equation, when $\partial_\theta \Sigma \equiv 0$, can be written as in Eq. (9) where $\partial_\theta Z(x)$ (11) can be obtained by deriving Eq. (10).

REFERENCES

- Capriotti, L. 2007. Least squares importance sampling for libor market models. *Wilmott Magazine* November: 100–107.
- Capriotti, L. 2008. Least squares importance sampling for monte carlo security pricing. *Quantitative Finance (in press)*.
- Giles, M. 2007. Working paper. Technical report, Mathematical Institute, University of Oxford, Oxford, United Kingdom.
- Glasserman, P. 2004. *Monte carlo methods in financial engineering*. New York: Springer.
- Harrison, J. M., and D. M. Kreps. 1979. Martingales and arbitrage in multiperiod securities markets. *Journal of Economic Theory* 20 (3): 381–408.
- Hull, J. C. 2002. *Options, futures and other derivatives*. New Jersey: Prentice Hall.
- Jaekel, P. 2002. *Monte carlo methods in finance*. London: Wiley Finance.
- Kallenberg, O. 1997. *Foundations of modern probability*. New York: Springer.

- Milstein, G., and M. Tretyakov. 2005. Numerical analysis of monte carlo evaluation of greeks by finite differences. *Journal of Computational Finance* 8:1–34.
- Schonbucher, P. 2003. *Credit derivatives pricing models: Models, pricing, implementation*. London: Wiley Finance.

AUTHOR BIOGRAPHY

LUCA CAPRIOTTI is a Vice President at Credit Suisse Group, Investment Banking Division, where he works in the Global Modelling and Analytics Group (GMAG). He was previously a researcher at the Kavli Institute for Theoretical Physics, Santa Barbara, California, working in the field of High Temperature Superconductivity and Quantum Monte Carlo methods for Condensed Matter systems. His current interests are in the field of Computational Finance, mainly focusing on efficient numerical techniques for Derivatives Pricing and Risk Management.