Technical article

Risk is the world's leading financial risk management magazine. **Risk**'s Cutting Edge articles are a showcase for the latest thinking and research into derivatives tools and techniques, with comprehensive coverage of the most recent advances in areas such as option pricing and hedging, market risk, credit risk, swaps and Monte Carlo methods.

www.risk.net/technical

Hedge your Monte Carlo

While the traditional Black-Scholes approach to option pricing is appealing on grounds of both elegance and tractability, the assumptions underlying it are usually violated in real markets. Here, Marc Potters, Jean-Philippe Bouchaud and Dragan Šestović propose an alternative Monte Carlo-based variance reduction approach to pricing and hedging

The Black-Scholes (1973) options model has two remarkable properties: one can find a "perfect" hedging strategy that eliminates risk entirely, and the option price does not depend on the average return of the underlying asset (see, eg, Wilmott, 1998). The latter property shows that the option price is not simply the (discounted) average of the future payout over the objective (or historical) probability distribution, as one would have expected. This is even more striking in the case of the Cox-Ross-Rubinstein binomial model, where the pricing measure is unrelated to the actual distribution of returns. The requirement of absence of arbitrage opportunities is actually equivalent to the existence of a "risk-neutral probability measure" (*a priori* distinct from the objective one), such that the price of a derivative is its (discounted) average payout, but where the average is performed over the risk-neutral distribution rather than over the objective distribution. It is thus a common belief that the knowledge of the "true" probability distribution of returns is useless information when pricing options. The credence is rather that the relevant risk-neutral distribution is somehow "guessed" by the market.

However, in most models of stock fluctuations, except for very special cases, risk in option trading cannot be eliminated, and strict arbitrage opportunities do not exist, whatever the price of the option. That risk cannot be eliminated is, furthermore, the fundamental reason for the very existence of option markets. It would thus be more satisfactory to have a theory of options where the objective stochastic process followed by the underlying asset was used to calculate the option price, the hedging strategy and the residual risk. It is clearly important to estimate the last of these for risk control purposes. A natural framework for this is the risk minimisation approach developed by several authors (Schweizer, 1996, Bouchaud & Potters, 1997, and Laurent & Pham, 1999), where the optimal trading strategy is determined such that the chosen measure of risk is minimised (for example, the variance of the wealth balance, although other choices are possible (Bouchaud & Potters, 1997, and Selmi & Bouchaud, 2001). The "theoretical" price is then obtained using a fair game argument. Interestingly, this framework allows one to recover exactly the Black-Scholes results when the objective probabilities are log-normal, and when the continuous time limit is taken (this is shown in detail in Bouchaud & Potters, 1997), or the Cox-Ross-Rubinstein results in the binomial case. In particular, the average trend completely disappears from the price and hedge.

The aim of this article is to present a very general Monte Carlo scheme based on this approach, which we call "optimal hedged Monte Carlo" (OHMC). The method, which has been inspired in part by the least-square method (LSM) of Longstaff & Schwartz (2001), shares with it the property that it can price a wide variety of exotic options, including those with path-dependent or early exercise features. On top of that, the OHMC has at least four major advantages over the standard Monte Carlo (SMC) scheme, where paths are generated with a weight consistent with the risk-neutral distribution:

- ☐ The OHMC method provides not only a numerical estimate of the price of the derivative, but also of the optimal hedge (which may be different from the Black-Scholes Δ-hedge for non-Gaussian statistics) and of the residual risk.
- ☐ The OHMC method leads to considerable variance reduction. This is related to the fact that the financial risk arising from the imperfect replica-

Implementation

- ☐ Generate discrete paths with objective probabilities, using either a model (log-Brownian, Garch, multi-fractal, etc) or using historical time series. Start all paths at time $t = -1$ (or before). This allows one to calculate both the price and the Greeks at time $t = 0$.
- ☐ Associate with each path a value equal to the final payout along the path.
- ☐ Write the value of the option (and the hedge) on the second-to-last node as a linear combination of functions of the current spot price. Find the coefficients of this combination by minimising the residual financial risk (a linear least-square problem).
- ☐ For each path, compare the above value of the option with early exercise and other path-dependent features of the option. Redefine accordingly the value of the option on that path.
- ☐ Repeat the previous two steps going backwards in time until the initial time step $t = 0$ is reached. ■

Least-square method of Longstaff & Schwartz

The least-square method (LSM) used in Longstaff & Schwartz (2001) differs from the optimal hedged Monte Carlo in the following ways:

- ☐ The paths must be generated using the risk-neutral probability distribution.
- ☐ The hedge is not used in the least-square fit equation (4).
- ☐ On subsequent iterations, the option price is kept as the discounted payout on that particular path (final payout or early exercise value).
- ☐ The final option price is given as the average payout of all paths. In this framework, the least-square fit only serves to find the early exercise points. Therefore, for European-style options, the LSM is identical to the standard Monte Carlo. ■

American-style options

The optimal hedged Monte Carlo (OHMC) method can be used to reduce the Monte Carlo error for all types of exotic options. We illustrate this point by showing how the method can be extended to price an American-style put option. To implement the early exercise condition, one can simply replace $\mathcal{C}_{k+1}(x_{k+1})$ in equation (4) by $\max(\mathcal{C}_{k+1}(x_{k+1}), x_s - x_{k+1})$, where x_s is the strike price. We have chosen a slightly different implementation, where we first find the early exercise point x^*_{k+1} and exercise all options for which $x_{k+1} < x^*_{k+1}$.

We have tested the method on a one-year American-style put option on a stock following a standard log-Brownian process. We follow the choice of parameters made in Longstaff & Schwartz (2001) to compare our results with theirs. The initial price and the strike are set to $x_0 = x_s = 40$, the volatility to 20% a year and the risk-free rate and the drift to 6%. As a benchmark price we use the value 2.314 (quoted in Longstaff & Schwartz) calculated using a very accurate finite difference method. We calculated the price within the OHMC using $N_{MC} = 500$ paths and $M = 8$ basis functions. To measure the accuracy of the method, we ran the Monte Carlo 500 times with different random seeds. The average price found was 2.302 with a standard deviation (around the true value 2.314) of 0.032.

We also used the least-square method (LSM) of Longstaff & Schwartz with the same parameters ($N_{MC} = 500$ paths and $M = 8$ basis functions). The average price within the LSM is found to be 2.423 with a standard deviation around the true value of 0.170, five times larger than for the above quoted 0.032 for the OHMC. These numbers are compatible with those found in Longstaff & Schwartz, where the error quoted is 0.01, ie, 17 times smaller but with 200 times more paths and 2.5 times more intermediate points.

Obviously, the same variance reduction would hold for other exotic options, as those discussed in Longstaff & Schwartz. A barrier option example on Microsoft is given in the text. ■

1. Option price as a function of underlying price for an OHMC simulation

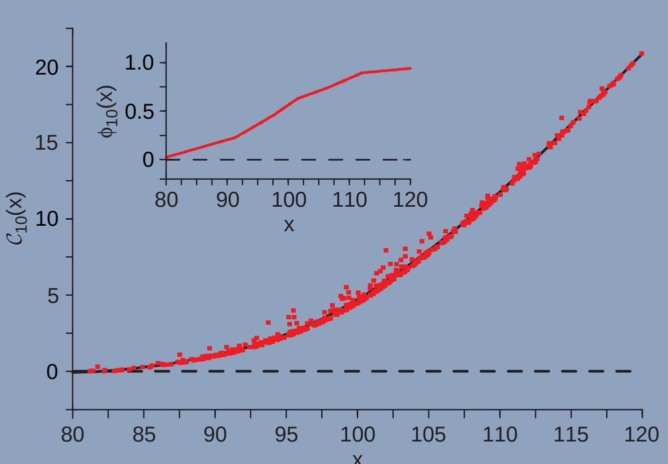

The full line corresponds to the fitted option price $\mathcal{C}_{10}(x)$ at the tenth hedging step ($k = 10$) of a simulation of length $N = 20$. Square symbols correspond to the option price on the next step, corrected by the hedge, equation (6), for individual Monte Carlo trajectories. The inset shows the hedge as a function of underlying price at the tenth step of the same simulation

tion of the option by the hedging strategy is directly related to the variance of the Monte Carlo simulation. When minimising the former by choosing the optimal strategy, we automatically reduce the latter. The standard deviation of our results are typically five to 10 times smaller than with the SMC, which means that for the same level of precision, the number of trajectories needed in the Monte Carlo is up to 100 times smaller.

☐ The method does not rely on the notion of risk-neutral measure and can be used with any model of the true dynamics of the underlying (even very complex ones), in particular those for which the risk-neutral measure is unknown and/or not uniquely defined.

☐ The OHMC method allows one to use purely historical data to price derivatives, short-circuiting the modelling of the underlying asset fluctuations. These fluctuations are known to be of a rather complex statistical nature, with fat-tailed distributions, long-range volatility correlations, negative return-volatility correlations, etc (Bouchaud & Potters, 1997, Guillaume *et al*, 1997, Mantegna & Stanley, 1999, and Muzy, Delour & Bacry, 2000). With the OHMC method, one can directly use the historical time series of the asset to generate the paths. The fact that a small number of paths is needed to reach good accuracy means that the length of the historical time series does not need to be very large.

Reduced variance Monte Carlo techniques for option pricing have been discussed in the literature (Clewlow & Caverhill, 1994), and bear some similarity with the present method. The general idea is to add some "control variates" to the observable one wants to average, which have by construction a zero average value but such that the resulting sum has smaller fluctuations. The profit and loss of some appropriate hedging strategy is an obvious candidate for such control variates, as was demonstrated in Clewlow & Caverhill (1994). However, the present method differs from the previous work in several key ways: first, the hedge used in Clewlow & Caverhill is an approximate hedge (for example, the Δ-hedge corresponding to a similar option for which an analytical formula is known), and not the optimal hedge for the option and underlying under consideration. Second, the idea of using the objective (historical) probability distribution is not discussed. Third, we couple the idea of hedged Monte Carlo with the versatile LSM of Schwartz & Longstaff (2001).

Basic principles of the method

Option pricing always requires working backwards in time. This is because the option price is exactly known at maturity, where it is equal to the payout. As with other schemes, we determine the option price by working step-by-step for maturity $t = N\tau$ to the present time $t = 0$, the unit of time τ being, for example, one day. The price of the underlying asset at time $k\tau$ is denoted as x_k and the price of the derivative is $\mathcal{C}_k$. We assume for simplicity that $\mathcal{C}_k$ only depends on x_k (and of course on k). However, the method can be generalised to account for a dependence of $\mathcal{C}_k$ on the volatility, interest rate, etc, or to price multi-dimensional options (such as interest rate derivatives). We therefore also introduce the hedge $\phi_k(x_k)$, which is the number of underlying assets in the portfolio at time k when the price is equal to x_k. Within a quadratic measure of risk, the price and the hedging strategy at time k are such that the variance of the wealth change between k and $k + 1$ is minimised. More precisely, we define the local "risk" $\mathcal{R}_k$ as:

$$\mathcal{R}_k = \left\langle \left(\mathcal{C}_{k+1}(x_{k+1}) - \mathcal{C}_k(x_k) + \phi_k(x_k)\left[x_k - x_{k+1}\right] \right)^2 \right\rangle_o \qquad (1)$$

where $\langle \ldots \rangle_o$ means that we average over the objective probability measure (and not the risk-neutral one). As shown in Bouchaud & Potters (1997), the functional minimisation of $\mathcal{R}_k$ with respect to both $\mathcal{C}_k(x_k)$ and $\phi_k(x_k)$ gives equations that allow one to determine the price and hedge, provided $\mathcal{C}_{k+1}$ is known. In the cases where the resulting minimal risk can be made to vanish (for example, within the Black-Scholes or Cox-Ross-Rubinstein models), all classical results of financial mathematics are reproduced. Note that we have not included interest rate effects in equation (1). When the interest rate r is non-zero, one should consider the following modified equation:

$$\mathcal{R}_k = \left\langle \left(e^{-\rho}\mathcal{C}_{k+1}\left(x_{k+1}\right) - \mathcal{C}_k\left(x_k\right) + \phi_k\left(x_k\right)\left[x_k - e^{-\rho}\, x_{k+1}\right]\right)^2 \right\rangle_o \quad (2)$$

where $\rho = r\tau$ is the interest rate over an elementary time step τ.

To implement this numerically, we parallel the LSM of Longstaff & Schwartz (2001), developed within a risk-neutral approach. We generate a set of N_{MC} Monte Carlo trajectories x_k^ℓ, where k is the time index and ℓ the trajectory index. We decompose the functions $\mathcal{C}_k$ and ϕ_k over a set of M appropriate basis functions $C_a(x)$ and $F_a(x)$:

$$C_k\left(x\right) = \sum_{a=1}^{M} \gamma_a^k C_a\left(x\right) \qquad \phi_k\left(x\right) = \sum_{a=1}^{M} \varphi_a^k F_a\left(x\right) \quad (3)$$

In other words, we solve the minimisation problem within the variational space spanned by the functions $C_a(x)$ and $F_a(x)$. This leads to a major simplification, since now we have a linear optimisation problem in terms of the coefficients γ_a^k, φ_a^k. These coefficients must be such that:

$$\sum_{\ell=1}^{N_{MC}} \left(e^{-\rho}\mathcal{C}_{k+1}\left(x_{k+1}^\ell\right) - \sum_{a=1}^{M} \gamma_a^k C_a\left(x_k^\ell\right) + \sum_{a=1}^{M} \varphi_a^k F_a\left(x_k^\ell\right)\left[x_k^\ell - e^{-\rho}\, x_{k+1}^\ell\right]\right)^2 \quad (4)$$

is minimised. Those N minimisation problems (one for each $k = 0, \ldots, N - 1$) are solved working backwards in time with $\mathcal{C}_N(x)$, the known final payout function.

Although in general the optimal strategy is not equal to the Black-Scholes Δ-hedge, the difference between the two is often small, and only leads to a second-order increase of the risk (Bouchaud & Potters, 1997). Therefore, one can choose to work within a smaller variational space and impose that:

$$\varphi_a^k \equiv \gamma_a^k \qquad F_a\left(x\right) \equiv \frac{dC_a\left(x\right)}{dx} \quad (5)$$

This will lead to exact results only for Gaussian processes, but reduces the computation cost by a factor of two. In practice, we have chosen these basis functions to be piece-wise linear for F_a and piece-wise quadratic for C_a, with breakpoints that adapt to the generated Monte Carlo paths.

We finish by noting that the OHMC method can be implemented using many different price processes, including models or data with fluctuating volatilities. In this case, one could let the function $\mathcal{C}_k$ depend not only on x_k but also on the value of some filtered past volatility σ_k. This would allow option prices to depend explicitly on the volatility and to calculate a vega. Vega hedging using market instruments could also be included to reduce the risk (and the Monte Carlo variance) further.

Numerical results for Black-Scholes

We have first checked our OHMC scheme when the paths are realisations of a (discretised) geometric random walk. We have priced an at-the-money three-month European-style option, on an asset with 30% annualised volatility and a drift equal to the risk-free rate, which we set at 5% a year. The number of time intervals N is chosen to be 20. The initial stock and strike price are $x_0 = x_s = 100$, and the corresponding Black-Scholes price is $\mathcal{C}_0^{BS} = 6.58$. The number of basis functions is $M = 8$. We run 500 simulations containing $N_{MC} = 500$ paths each, for which we extract the average price and standard deviation on the price. An example of the result of linear regression is plotted in figure 1. Each data point corresponds to one trajectory of the Monte Carlo at one instant of time k, and represents the quantity:

$$e^{-\rho}\mathcal{C}_{k+1}\left(x_{k+1}\right) + \phi_k\left(x_k\right)\left[x_k - e^{-\rho}\, x_{k+1}\right] \quad (6)$$

as a function of x_k. The full line represents the result of the least-squared fit, from which we obtain $\mathcal{C}_k(x_k)$. We show in the inset the corresponding hedge ϕ_k, which was constrained in this case to be the Δ-hedge.

We obtain the following numerical results. For the SMC (unhedged) scheme, we obtain as an average over the 500 simulations $\mathcal{C}_0^{RN} = 6.68$ with a standard deviation of 0.44. For the OHMC, we obtain $\mathcal{C}_0^{H} = 6.55$ with a standard deviation of 0.06, seven times smaller than with the SMC scheme. This variance reduction is illustrated in figure 2, where we show the histogram of the 500 different Monte Carlo results both for the unhedged case (full bars) and for the hedged case (dotted bars). A similar variance reduction is reported in Clewlow & Caverhill (1994).

Now we set the drift to 30% a year. The Black-Scholes price, as is well known, is unchanged. A naive unhedged Monte Carlo scheme with objective probabilities gives a completely wrong price of 10.72, 60% higher than the correct price, with a standard deviation of 0.56. On the other hand,

2. Histogram of the option price as obtained from 500 MC simulations with different seeds

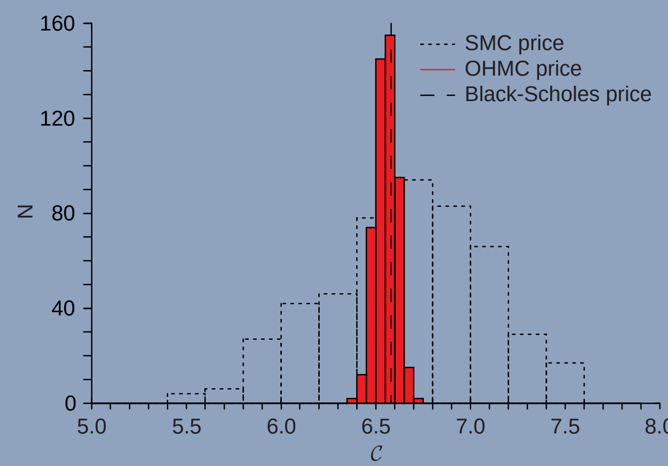

The dotted histogram corresponds to the SMC and the full histogram to the OHMC. The dotted line indicates the exact Black-Scholes price. Note that on average both methods give the correct price, but that the OHMC has an error that is more than seven times smaller than that of the SMC

3. Smile curve for a purely historical OHMC of a one-month option on Microsoft (volatility as a function of strike price)

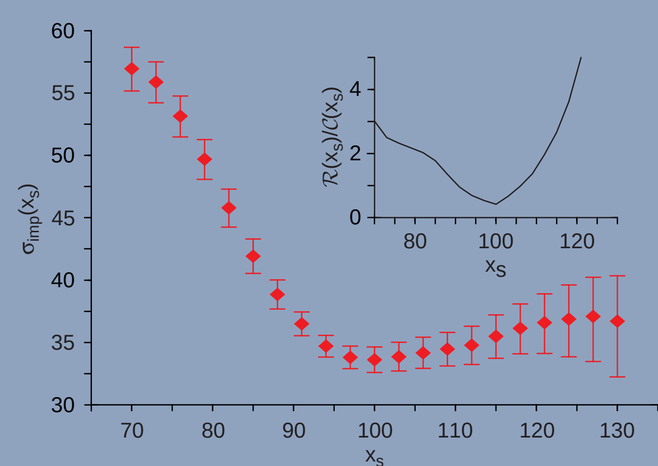

The error bars are estimated from the residual risk. The inset shows this residual risk as a function of strike normalised by the "time-value" of the option (ie, by the call or put price, whichever is out-of-the-money)

A. Results for a down-and-out call option on Microsoft

Maturity/strike	Implied Black-Scholes	OHMC
Two weeks/100	\$2.45	\$2.39
Two weeks/105	\$0.86	\$0.84
One month/100	\$3.10	\$3.08
One month/105	\$1.61	\$1.61

The barrier is at \$95. The spot price is normalised to \$100

the OHMC indeed produces the correct price (6.52), with a standard deviation of 0.06. The SMC scheme in this case simply amounts to setting "by hand" the drift to the risk-free rate, and therefore obviously gives back the above figures.

Therefore, we have checked that in the case of a geometric random walk, the OHMC indeed gets rid of the drift and reproduces the usual Black-Scholes results, as it should. This allows us to confidently extend the method to other types of option and other random processes. We have not investigated in depth the optimal values to be given to the parameters M and N_{MC}, or the choice of the basis functions that minimise the computation cost for a given accuracy. These are implementation issues that are beyond the scope of this article.

Purely historical option pricing

We now turn to the idea of a purely historical OHMC pricing scheme. We price a one-month (21 business days) option on Microsoft, hedged daily, with zero interest rates. We used 2,000 paths of length 21 days, obtained from the time series of Microsoft in the period May 1992 to May 2000. The initial price is always normalised to 100. We use a set of $M = 10$ basis functions, and stay with the simple Δ-hedge. From our numerically determined option prices, we extract an implied Black-Scholes volatility by inverting the Black-Scholes formula and plot it as a function of the strike, in order to construct an implied volatility smile. The result is shown in figure 3. Since we now only perform a single Monte Carlo simulation, the error bars shown are obtained from the residual risk of the optimally hedged options. The residual risk itself, divided by the call or the put option price (respectively for out-of-the-money and in-the-money call options), is given in the inset. We find that the residual risk is around 42% of the option premium at-the-money, and rapidly reaches 100% when one goes out-of-the-money. These risk numbers are comparable to those obtained on other options of similar maturity (see Bouchaud & Potters, 1997), and are much larger than the residual risk that one would get from discrete time hedging effects in a Black-Scholes world.

The smile that we obtain has a shape quite typical of those observed on option markets. However, it should be emphasised that we have neglected the possible dependence of the option price on the local value of the volatility.

It is interesting to price some exotic options within the same framework. In table A, we compare, for example, the price of a down-and-out call on Microsoft using the OHMC method and using the standard Black-Scholes formula for barrier options with the implied volatility corresponding to the same strike, as shown in figure 3. As could be expected, the OHMC price is lower than the Black-Scholes price, reflecting the presence of jumps in the underlying that increases the probability to hit the barrier. Interestingly, however, the difference becomes small as the maturity increases.

Conclusion and prospects

We have presented what we believe to be a useful Monte Carlo scheme, which closely follows the actual history of a trader hedged portfolio. The inclusion of the optimal hedging strategy allows one to reduce the financial risk associated with option trading and, for the same reason, the variance of our OHMC scheme as compared with the standard Monte Carlo schemes. The explicit accounting of the hedging cost naturally converts the objective probability into the "risk-neutral" one and allows one to recover all classical results of mathematical finance when markets are complete. This allows a consistent use of purely historical time series to price derivatives and obtain their residual risk. We believe that there are many extensions and applications of the scheme, for example, to price exotic options or interest rate derivatives using faithful historical models, and market hedging instruments. With some modifications and extra numerical cost, the method presented here could be used to deal with transaction costs, or with non-quadratic risk measures (value-at-risk hedging). Finally, it is interesting to adapt the present method to extract from market option prices implied parameters that can be used to price exotic options and for risk control. Work on this is in progress. ■

Marc Potters, Jean-Phillippe Bouchaud and Dragan Šestović are all at Science & Finance, the research division of Capital Fund Management. Jean-Phillippe Bouchaud is also at the Service de Physique de l'Etat Condensé, CEA Saclay. **We thank Jean-Pierre Aguilar, Andrew Matacz and Martin Bazant for interesting discussions**

Comments on this article can be posted on the technical discussion forum on the Risk website at http://www.risk.net

REFERENCES

Black F and M Scholes, 1973
The pricing of options and corporate liabilities
Journal of Political Economy 81, pages 637–659

Bouchaud J-P and M Potters, 1997
Theory of financial risks (in French)
Aléa-Saclay, Eyrolles, Paris (in English, Cambridge University Press, 2000)

Bouchaud J-P and M Potters, 1999
Back to basics: historical option pricing revisited
Philosophical Transactions: Mathematical, Physical & Engineering Sciences 357(1,758), pages 2,019–2,028

Clewlow L and A Caverhill, 1994
On the simulation of contingent claims/Quicker on the curves
Journal of Derivatives 2, pages 66–74/*Risk* May, pages 63–65

Guillaume D, M Dacorogna, R Davé, U Müller, R Olsen and O Pictet, 1997
From the bird's eye to the microscope: a survey of new stylized facts of the intra-daily foreign exchange markets
Finance and Stochastics 1, pages 95–129

Laurent J-P and H Pham, 1999
Dynamic programming and mean-variance hedging
Finance and Stochastics 3, pages 83–110

Longstaff F and E Schwartz, 2001
Valuing American options by simulation: a simple least-squares approach
Review of Financial Studies 14, pages 113–147

Mantegna R and H Stanley, 1999
An introduction to econophysics
Cambridge University Press

Muzy J-F, J Delour and E Bacry, 2000
Modelling fluctuations of financial time series: from cascade process to stochastic volatility model
European Physical Journal B 17, pages 537–548

Plerou V, P Gopikrishnan, L Amaral, M Meyer and H Stanley, 1999
Scaling of the distribution of price fluctuations of individual companies
Physical Review E 60, pages 6,519–6,529

Schweizer M, 1996
Approximation pricing and the variance optimal martingale measure
Annals of Probability 24, pages 206–236

Selmi F and J-P Bouchaud, 2001
Hedging large risks reduces transaction costs
Forthcoming in Wilmott Magazine

Wilmott P, 1998
Derivatives, the theory and practice of financial engineering
John Wiley & Sons