

A Monte Carlo method for exponential hedging of contingent claims

M. R. Grasselli* and T. R. Hurd*
Dept. of Mathematics and Statistics
McMaster University
Hamilton ON L8S 4K1
Canada

July 8, 2004

Abstract

Utility based methods provide a very general theoretically consistent approach to pricing and hedging of securities in incomplete financial markets. Solving problems in the utility based framework typically involves dynamic programming, which in practise can be difficult to implement. This article presents a Monte Carlo approach to optimal portfolio problems for which the dynamic programming is based on the exponential utility function $U(x) = -\exp(-x)$. The algorithm, inspired by the Longstaff-Schwartz approach to pricing American options by Monte Carlo simulation, involves learning the optimal portfolio selection strategy on simulated Monte Carlo data. It shares with the LS framework intuitivity, simplicity and flexibility.

*Research supported by the Natural Sciences and Engineering Research Council of Canada and Mathematics of Information Technology and Complex Systems, Canada

1 Introduction

As realized in the pioneering work of Black, Scholes, Merton and others, financial assets in complete markets can be priced uniquely by construction of replicating portfolios and application of the no arbitrage principle. This conceptual framework forms the basis of much of the currently used methodology for financial engineering. In recent years, however, finance practitioners have been increasingly led by competitive pressures to the use of more general incomplete market models, such as those driven by noise with stochastic volatility, jumps or general Lévy processes. In incomplete markets, matters are much more complicated, and the pricing and hedging of financial assets depends on the risk preferences of the investor.

Utility based pricing and hedging are extensions growing naturally out of portfolio optimization, and much work is now in progress to place these methods in the broadest context and to explore their various ramifications. This framework leads to new concepts, notably the Davis price [8] and the indifference price of the contingent claim [18]. This general theory is naturally applicable in areas where the complete market theory appears inappropriate, such as real options [16], insurance [27] and the general valuation of non-traded assets [15, 28]. In the insurance context, the indifference price can be thought of as the reservation price of the claim, that is the amount the insurer should set aside to deal with its future liability.

The basic problem is that of a rational agent who seeks to find their optimal hedging portfolio when they have sold (or bought) a contingent claim B which matures at time T . In the language of utility theory, the agent tries to solve the problem

$$u(x) = \sup_{H \in \mathcal{A}} E \left[U \left(X_T^{x,H} - B \right) \right], \quad (1)$$

for a concave, strictly increasing, differentiable utility function $U : \mathbb{R} \rightarrow \mathbb{R} \cup \{-\infty\}$. Here $X_T^{x,H}$ denotes the agent's discounted wealth at time T achieved by adopting a portfolio strategy H starting with initial wealth x . In the absence of the claim B , we see that (1) is reduced to Merton's optimal investment problem, for which key results were obtained in [19, 20, 21]. In general, B is the \mathcal{F}_T -measurable random variable corresponding to the discounted liability to

be faced at time T .

If $S_t = (S_t^1, \dots, S_t^d)$ is the \mathbb{R}^d -valued process which describes the discounted prices of traded assets, and $H_t = (H_t^1, \dots, H_t^d)$ is an \mathbb{R}^d -valued process representing the agent's asset allocations, then

$$X_T^{x,H} = x + (H \cdot S)_0^T.$$

The change in wealth over a period $[s, t]$ achieved by the *self financing* portfolio is the stochastic integral $(H \cdot S)_s^t := \sum_{i=1}^d \int_s^t H_u^i dS_u^i$ defined for example in [25], where the integrators S_t are taken to be *càdlàg* semimartingales on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \in [0, T]}, P)$. In financial terms this framework offers the flexibility of including discrete time and jump processes, rather than just continuous diffusions. The integrands H_t^i are assumed to be predictable S -integrable processes, which has the financial interpretation of investment decisions based on information available immediately before the investment date. They include as special cases portfolios which are piecewise constant but jump at discrete times, which correspond to the realistic situations when continuous reallocation of funds is impossible to implement. The domain of optimization is some convex set \mathcal{A} of admissible self-financing portfolios, which must be restrictive enough to rule out trading strategies for which the agent's wealth assumes arbitrarily negative values (such as “doubling strategies”) while still allowing for an optimal strategy to exist. For instance [14], one might take it to be the set for which the wealth process $X_t^{x,H}$ is uniformly bounded from below.

A general approach to solve (1) is via convex duality, by means of which the utility maximization problem over the appropriate class of admissible portfolios (the “primal problem”) is related to a minimization problem over a suitable domain in the set of measures on Ω (the “dual problem”). Let us denote the set of equivalent local martingale measures for the price process S by $\mathcal{M}^e(S)$ and assume that $\mathcal{M}^e(S) \neq \emptyset$ (which, for markets where the price process S is locally bounded, is equivalent to “No Free Lunch with Vanishing Risk”, a slightly restrictive notion of “No Arbitrage” [10, sections 2 and 3]). For bounded claims and under economically acceptable technical conditions on the class of utility functions (such as “reasonable asymptotic

elasticity”) and admissible portfolios, one has that both the primal and dual problems have unique optimizers $\widehat{X}(x)$ and $\widehat{Q}(y)$ satisfying the fundamental equation

$$U'(\widehat{X}(x) - B) = y \frac{d\widehat{Q}(y)}{dP}, \quad (2)$$

where x and y are related by $u'(x) = y$ [7, 9, 24].

Practical implementation of incomplete market models based on these new theoretical developments requires the development of efficient numerical methods. Three distinct approaches can be considered and ultimately all three are needed for a complete understanding of implementation issues. One approach, adopted for example in [13] is the numerical solution of general Hamilton–Jacobi–Bellman equations, which are the partial differential equations derived from stochastic control theory. A second approach could be broadly classified as “state space discretization”, by which we mean Markov chain or lattice based methods [23]. A third broad approach can be called Monte Carlo or random simulation methods. It is this third approach we attempt to realize in the present paper.

To our knowledge, Monte Carlo methods, although widely used for pricing derivatives [4, 12], have not been extensively used for optimal portfolio theory. Some works related to this in the context of complete markets are [11] and [6]. Our proposed application of Monte Carlo is intrinsically more difficult than for example its use in the pricing of American style options, a problem which has only quite recently been efficiently implemented with the least squares algorithm of [22]. Despite these difficulties, which we will see quite clearly in this paper, Monte Carlo methods have a great asset in being very simple and intuitive. By implementing such methods, we can gain key intuition and understanding which may be difficult to learn from the abstract theory.

The remaining of the paper is organized as follows. Section 2 describes optimal hedging strategies in discrete time, in particular the concepts of certainty equivalent value, indifference price and the Davis price. The main innovation of the paper is the exponential utility algorithm given in section 3. It is a Monte Carlo method for learning the optimal trading strategy in a

Markovian market for the class of discrete time hedging problems introduced in section 2. This algorithm is inspired by, but is quite different from, the least-squares algorithm of Longstaff and Schwartz for pricing American options. Interestingly, our method works well only for the exponential utility: our suggested extension to general utilities is much less efficient. Towards the end of this section we discuss the systematic sources of error arising from the algorithm.

In section 4 we present two preliminary numerical implementations of our algorithm. The first one is for the exactly solvable one-dimensional geometric Brownian motion model, where we compare the results obtained from our Monte Carlo simulations with those arising from the explicit theoretical formulas for both the Merton and hedging portfolios. In order to analyze the performance of the algorithm in detail, we compute several different risk measures of the profit/loss empirical distribution realized at terminal time for both the learned and theoretical trading strategies. Our second example is a two factor stochastic volatility model, inequivalent to but as rich as the Heston model [17], for which tractable expressions for indifference prices and hedging portfolios for *pure volatility claims* were obtained using an alternative method in [13]. We complement the results of [13] by computing the indifference prices and hedging portfolios for general claims written on the traded asset and the volatility factor. In our concluding section 5, we discuss the various advantages and drawbacks we observe in the method.

2 Discrete time hedging

We consider discrete time hedgings, that is, in addition to the criteria of admissibility mentioned in the introduction, we restrict the class \mathcal{A} of admissible portfolios to processes of the form

$$H_t = \sum_{k=1}^K H_k \mathbf{1}_{(t_{k-1}, t_k]}(t), \quad (3)$$

where each H_k is an \mathbb{R}^d -valued $\mathcal{F}_{k-1} := \mathcal{F}_{t_{k-1}}$ random variable (we write $H_k \in \mathcal{F}_{k-1}$). That is, based on the information available at time t_{k-1} the investor chooses the portfolio allocation H_k to be held until the next reallocation time t_k . We take the discrete time partition of the interval

$[0, T]$ to be of the form

$$t_0 = 0 < t_1 = \frac{T}{K} < \dots < t_k = \frac{kT}{K} \dots < t_K = T$$

and use the notation $S_j := S_{t_j}$ for discrete time stochastic processes. The discounted wealth for self-financing portfolios is

$$X_j = x + (H \cdot S)_0^j, \quad (4)$$

with the notation $(H \cdot S)_k^j := \sum_{i=1}^d \sum_{l=k+1}^j H_l^i \Delta S_l^i$ and $\Delta S_k := S_k - S_{k-1}$.

To better understand the optimal selection problem it is useful to formulate a dynamical version of it. Let us write H_t^B for the optimal solution to the static primal problem (1). For any intermediate time $t_k \in [0, T]$ and $x \in \text{dom}(U)$, we can write

$$u(x) = \sup_{H \in \mathcal{A}(0, t_k]} E \left[\text{ess sup}_{H \in \mathcal{A}(t_k, T]} E_k[U(x + (H \cdot S)_0^k + (H \cdot S)_k^K - B)] \right], \quad (5)$$

where we use $E_k[\cdot]$ to denote the conditional expectation with respect to \mathcal{F}_k . This leads us to the study of the *conditional* problem

$$u_k(\mathbf{w}) = \text{ess sup}_{H \in \mathcal{A}(t_k, T]} E_k[U(\mathbf{w} + (H \cdot S)_k^K - B)], \quad (6)$$

where $\mathbf{w} \in \mathcal{F}_k$ represents the stochastic wealth accumulated up to time t_k . If we trade according to H_t^B up to time t , that is, if $\mathbf{w} = x + (H^B \cdot S)_0^k$, then we must have

$$u_k(\mathbf{w}) = E_k[U(\mathbf{w} + (H^B \cdot S)_k^K - B)], \quad (7)$$

for the restriction H_u^B , $t_k \leq u \leq T$. In other words, the optimal portfolio H_t^B is also conditionally optimal. This is a special instance of the dynamic programming principle, which for this discrete time stochastic control problem has the form of K subproblems

$$u_{k-1}(\mathbf{w}) = \text{ess sup}_{H_k \in \mathcal{F}_{k-1}} E_{k-1}[u_k(\mathbf{w} + (H \cdot S)_{k-1}^k)], \quad k = K, K-1, \dots, 1, \quad (8)$$

subject to the terminal condition $u_K(x) = U(x - B)$.

There is a useful way to view the value function $u_k(\mathbf{w})$. Since the utility function is invertible on its domain, for each (\mathbf{w}, t_k) we can define the *certainty equivalent value* of the claim B at time t as the random variable $c_k^B(\mathbf{w})$ satisfying

$$U(\mathbf{w} - c_k^B(\mathbf{w})) = E_k[U(\mathbf{w} + (H^B \cdot S)_k^K - B)]. \quad (9)$$

That is, the certain utility achieved by investing the amount $\mathbf{w} - c_k^B(\mathbf{w})$ in the risk free account equals the expected utility of the terminal wealth $\mathbf{w} + (H^B \cdot S)_k^K - B$ of the optimal hedging portfolio. From (8), we see that the certainty equivalent process $c_k^B(\mathbf{w})$ satisfies

$$U(\mathbf{w} - c_{k-1}^B(\mathbf{w})) = \operatorname{ess\,sup}_{H_k \in \mathcal{F}_{k-1}} E_{k-1}[U(\mathbf{w} + H_k \Delta S_k - c_k^B(\mathbf{w} + H_k \Delta S_k))] \quad (10)$$

with $c_K^B(\mathbf{w})$ taken equal to the terminal discounted claim B . Therefore, $c_k^B(\mathbf{w})$ represents a wealth dependent effective value of the claim B at time t_k .

Following [18] (according to [2]), a clear interpretation of the certainty equivalent can be given by considering an investor who, holding wealth \mathbf{w} at time t_k , must decide the minimum amount π to charge when selling a claim B . If he sells the claim for π and hedges optimally against the claim he will achieve a maximal expected utility

$$\operatorname{ess\,sup}_{H \in \mathcal{A}(t_k, T]} E_k[U(\mathbf{w} + (H \cdot S)_k^K - B)] = U(\mathbf{w} + \pi - c_k^B(\mathbf{w} + \pi))$$

If, however, the investor does not sell the claim and invests optimally according to the solution for Merton's problem, he achieves

$$\operatorname{ess\,sup}_{H \in \mathcal{A}(t_k, T]} E_k[U(\mathbf{w} + (H \cdot S)_k^K)] = U(\mathbf{w} - c_k^0(\mathbf{w})).$$

The *seller's indifference price* of the claim B at time t_k for wealth \mathbf{w} is the value $\pi = \pi_t^B(\mathbf{w})$

which makes these equal, that is, it is the solution (if it exists) for

$$\pi_k^B(\mathbf{w}) = c_k^B(\mathbf{w} + \pi_k^B(\mathbf{w})) - c_k^0(\mathbf{w}). \quad (11)$$

To obtain the correct notion of a “buyer’s price” we need to consider the reverse claim $-B$ in (1): the *buyer’s indifference price* for the claim B at time t_k and wealth \mathbf{w} is given by

$$\tilde{\pi}_k^B(\mathbf{w}) = -\pi_k^{-B}(\mathbf{w}). \quad (12)$$

If, for each ε , we let $\pi_k^{\varepsilon B}(\mathbf{w})$ denote the indifference price of the claim εB , then the *Davis price* [8] of the claim B is defined to be

$$\pi_k^{Davis}(\mathbf{w}) = \left. \frac{d\pi_k^{\varepsilon B}(\mathbf{w})}{d\varepsilon} \right|_{\varepsilon=0}. \quad (13)$$

By differentiating the identity

$$U(\mathbf{w} - c_k^{\varepsilon B}(\mathbf{w})) = E_k[U(\mathbf{w} + (H^{\varepsilon B} \cdot S)_k^K - \varepsilon B)]$$

at $\varepsilon = 0$ it is straightforward to see that the Davis price of B is given by the expectation pricing

$$\pi_k^{Davis}(\mathbf{w}) = E_k^{\widehat{Q}(y)}[B]. \quad (14)$$

where $\widehat{Q}(y)$ stands for the optimal solution to the dual to Merton’s problem. We remark that the indifference price, being intrinsically nonlinear, does not in general satisfy useful criteria such as put-call parity. The Davis price, on the other hand, does.

To conclude this section, we observe what happens if there exists a portfolio \mathcal{H}^B of the form (3) which replicates the claim B in the sense that $B = B_0 + (\mathcal{H}^B \cdot S)_0^T$ for some constant B_0 (this is true for any claim in complete markets). Then we can write (1) in the form of the Merton

problem

$$\sup_{H \in \mathcal{A}} E [U(x - B_0 + ((H - \mathcal{H}^B) \cdot S)_0^T)].$$

If $\widehat{H}(x - B_0)$ is the optimal portfolio for the Merton problem starting with wealth $x - B_0$, it follows that the optimal portfolio for the hedging problem for the claim B starting with wealth x will be given by

$$H^B(x) = \widehat{H}(x - B_0) + \mathcal{H}^B, \quad (15)$$

and the indifference price for the claim B can be expressed as

$$\pi_k^B = B_0 + (\mathcal{H}^B \cdot S)_0^k. \quad (16)$$

3 The exponential utility allocation algorithm

In this section we introduce a Monte Carlo method for learning the optimal trading strategy (8) for the discrete time problems of the previous section. We want an algorithm which will generate an approximate trading rule, based on a data set $\{(S_k^i, Y_k^i)\}_{i=1, \dots, N; k=1, \dots, K}$ where $(S_k^i, Y_k^i) \in \mathbb{R}^n$ denotes the state of the i th sample path at time $t_k = kT/K$. We fix the initial values to be $(S_0^i, Y_0^i) = (S_0, Y_0)$. In what follows, we will be largely concerned with an exponential utility function and with markets and claims satisfying Markovian conditions. More explicitly, we assume the following:

Assumption 1 *The market is Markovian and its state variables $(S^1, \dots, S^d, Y^1, \dots, Y^{n-d})$ lie in a finite dimensional state space $\mathcal{S} \subset \mathbb{R}^n$.*

Assumption 2 *The contingent claim is taken to be of the form $B_T = B(S_T, Y_T)$ for a bounded Borel function $B : \mathcal{S} \rightarrow \mathbb{R}$.*

Assumption 3 *The utility function has the form $U(x) = -\frac{e^{-\gamma x}}{\gamma}$, $\gamma > 0$.*

In these assumptions, we interpret S as discounted asset prices as before and the additional variables Y as values of nontraded quantities such as stochastic volatilities which may or may

not be observed directly.

As for the choice of an exponential utility, consider the discrete time problem (8) for a general utility function U in this Markovian setting. The optimal portfolio $(H^B)_{k+1}^i \in \mathbb{R}^d$ should be selected as a function of (X_k^i, S_k^i, Y_k^i) where X_k^i is the wealth held at the point (i, k) . We can see a basic difficulty with a Monte Carlo approach: to “learn” the function for H_{k+1}^B from the data $\{(S, Y)\}$ will require being able to fill in the optimal wealth from time t_0 to t_k . One way to do this would be to perform an additional simulation of the pair (X_k^i, X_{k+1}^i) at each point (i, k) : this would however result in a drastic drop in efficiency. In contrast, for the special case of an exponential utility, a look at (6) shows that $u_k(\mathbf{w})$ factorizes as

$$u_k(\mathbf{w}) = -\frac{e^{-\gamma \mathbf{w}}}{\gamma} \operatorname{ess\,inf}_{H \in \mathcal{A}(t_k, T]} E_k \left[e^{-\gamma(H \cdot S)_k^K + \gamma B} \right] =: -\frac{e^{-\gamma \mathbf{w}}}{\gamma} v_k, \quad (17)$$

so that the reduced value function v_k , as well as the certainty equivalent c_k^B and the optimal hedging portfolio H_k^B are all wealth independent processes. In the Markovian setting this means that they have the form

$$v_k = v_k(S_k, Y_k) \quad (18)$$

$$c_k^B = c_k(S_k, Y_k) \quad (19)$$

$$H_{k+1}^B = h_{k+1}(S_k, Y_k) \quad (20)$$

for (deterministic) Borel scalar functions $\{v_k, c_k\}_{k=0}^{K-1}$ and \mathbb{R}^d -valued functions $\{h_{k+1}\}_{k=0}^{K-1}$ on the state space \mathcal{S} . For this reason our algorithm is at this point restricted to exponential utility functions, and we take $\gamma = 1$ for simplicity.

3.1 The algorithm

First observe that, if H_K^B is the optimizer of (8) for $k = K$ then for any other random variable $H_K \in \mathcal{F}_{K-1}$ we have

$$\begin{aligned} E[U(x + (H_K^B \cdot \Delta S_K) - B)] &= E[E_{K-1}[U(x + (H_K^B \cdot \Delta S_K) - B)]] \\ &\geq E[E_{K-1}[U(x + (H_K \cdot \Delta S_K) - B)]] \\ &= E[U(x + (H_K \cdot \Delta S_K) - B)], \end{aligned}$$

so that H_K^B is also the optimizer for

$$\sup_{H_K \in \mathcal{F}_{K-1}} E[U(x + (H_K \cdot \Delta S_K) - B)], \quad (21)$$

and similarly for the other time steps. With this in mind, we proceed to describe our algorithm as follows.

1. *Step $k = K$:* The final optimal allocation H_K^B is the \mathbb{R}^d -valued \mathcal{F}_{K-1} -random variable which solves

$$\min_{H_K \in \mathcal{F}_{K-1}} E[\exp(-H_K \cdot \Delta S_K + B)]. \quad (22)$$

Since the solution is known to be given by $H_K^B = h_K(S_{K-1}, Y_{K-1})$ for some deterministic function $h_K \in \mathcal{B}(\mathcal{S})$ (the set of Borel functions on \mathcal{S}), we write this as

$$\min_{h_K \in \mathcal{B}(\mathcal{S})} E[\exp(-h_K(S_{K-1}, Y_{K-1}) \cdot \Delta S_K + B)]. \quad (23)$$

One cannot hope to determine h_K on only a finite set of data: we therefore pick an R -dimensional subspace $\mathcal{R}(\mathcal{S}) \subset \mathcal{B}(\mathcal{S})$ of functions on \mathcal{S} and “learn” a suboptimal solution

$$\arg \min_{h_K \in \mathcal{R}(\mathcal{S})} E[\exp(-h_K(S_{K-1}, Y_{K-1}) \cdot \Delta S_K + B)].$$

By the central limit theorem, the expectation above can be approximated by the finite

sample estimate

$$\Psi_K(h_K) = \frac{1}{N} \sum_{i=1}^N \exp(-h_K(S_{K-1}^i, Y_{K-1}^i) \cdot \Delta S_K^i + B(S_K^i, Y_K^i)). \quad (24)$$

This leads to our estimator $h_K^{\mathcal{R}}$ based on $\{S_k^i, Y_k^i\}$ and the choice of subspace \mathcal{R} :

$$h_K^{\mathcal{R}} := \arg \min_{h_K \in \mathcal{R}(S)} \Psi_K(h_K). \quad (25)$$

2. *Inductive step for $k = K - 1, \dots, 2$:* The estimator $h_k^{\mathcal{R}}$ of the optimal rule h_k , for the intermediate time steps $2 \leq k < K - 1$ is determined inductively given the estimators $h_{k+1}^{\mathcal{R}}, \dots, h_K^{\mathcal{R}}$. It is defined to be

$$h_k^{\mathcal{R}} := \arg \min_{h_k \in \mathcal{R}(S)} \Psi_k(h_k; h_{k+1}^{\mathcal{R}}, \dots, h_K^{\mathcal{R}}) \quad (26)$$

where

$$\Psi_k(h_k) = \frac{1}{N} \sum_{i=1}^N \exp(-h_k(S_k^i, Y_k^i) \cdot \Delta S_{k+1}^i + c_k^i(h_{k+1}^{\mathcal{R}}, \dots, h_K^{\mathcal{R}}, S_k^i, Y_k^i)), \quad (27)$$

with

$$c_k^i(h_{k+1}^{\mathcal{R}}, \dots, h_K^{\mathcal{R}}) = B(S_K^i, Y_K^i) - \sum_{j=k+1}^K h_j^{\mathcal{R}}(S_{j-1}^i, Y_{j-1}^i) \cdot \Delta S_j^i. \quad (28)$$

3. *Final step $k = 1$:* This step is degenerate since the initial values (S_0, Y_0) are constant over the sample. Therefore we determine the optimal constant vector $h_1 \in \mathbb{R}^d$ by solving

$$h_1 := \arg \min_{h \in \mathbb{R}^d} \Psi_1(h; h_2^{\mathcal{R}}, \dots, h_K^{\mathcal{R}}). \quad (29)$$

To summarize, the algorithm above learns a collection of functions of the form

$$(h_1, h_2^{\mathcal{R}}, \dots, h_K^{\mathcal{R}}) \in \mathbb{R}^d \times \mathcal{R}(S)^{K-1}$$

from the Monte Carlo simulation. This collection defines a suboptimal allocation strategy for the exponential hedging problem. The optimal value

$$\Psi_1(h_1; h_2^{\mathcal{R}}, \dots, h_K^{\mathcal{R}}) = \frac{1}{N} \sum_{i=1}^N \exp \left(-h_1(S_0, Y_0) - \sum_{j=2}^K h_j^{\mathcal{R}}(S_{j-1}^i, Y_{j-1}^i) \cdot \Delta S_j^i + B(S_K^i, Y_K^i) \right), \quad (30)$$

is an estimate of the quantity $\exp(c_0^B)$, where c_0^B is the certainty equivalent value of the claim B at time $t = 0$. Finally, the indifference price of the claim is approximated by $\log(\Psi_1/\tilde{\Psi}_1)$ where $\Psi_1, \tilde{\Psi}_1$ are given by (30) with the claims B and 0 respectively.

3.2 Systematic errors

It is important to identify two distinct systematic sources of error in the algorithm. The first, which we call *approximation one* (following [5]), is in focusing on suboptimal solutions $h_k^{\mathcal{R}}$ which lie in a specified subspace $\mathcal{R}(\mathcal{S})$ of the full space $\mathcal{B}(\mathcal{S})$. From a pragmatic perspective, we need to select a set of R basis functions f_1, \dots, f_R for $\mathcal{R}(\mathcal{S})$ which does a good job of representing the true optimal function over the values of state space covered by the Monte Carlo simulation. Naively, one might expect to need to choose R exponentially related to the dimension of \mathcal{S} ; experience seems to suggest that far fewer functions are needed for higher dimension problems. For a discussion of this type of question in the context of the Longstaff-Schwartz (LS) method for American options, see [22] and [5]. Observe that the requirements of our algorithm are much more stringent than for the American option problem, since the strategy to be learned is not simply “to exercise or not to exercise”, but must select a high dimensional vector at each point (i, k) in the simulation. Having said this, we take the point of view that the careful selection of a subspace $\mathcal{R}(\mathcal{S})$ might lead to good performance of the algorithm. Furthermore, our experiments show that the sensitivity to changes in $\mathcal{R}(\mathcal{S})$ of quantities such as indifference prices are much less than that of quantities such as hedge allocations.

The second source of error, *approximation two*, is the finite N approximation. We can in principle estimate this error in terms of the basic model parameters. A detailed analysis for the one period problem is offered in [1]. The heuristics of that argument is the following. For

$k \leq K$, denote by $I_k(h_k)$ the true expectation being approximated by the Monte Carlo sample sum $\Psi_k(h_k)$, that is

$$I_k(h_k) := E[\exp(-h_k(S_{k-1}, Y_{k-1}) \cdot \Delta S_k + c_k(S_k, Y_k)]. \quad (31)$$

By the central limit theorem, for a given confidence level $1 - \alpha$, $\alpha \ll 1$, there exist constants C_1, C_2 such that, with probability $1 - \alpha$, we have

$$|\Psi_k(h_k) - I_k(h_k)| \leq \frac{C_1}{\sqrt{N}}, \quad \|\nabla \Psi_k(h_k) - \nabla I_k(h_k)\| \leq \frac{C_2}{\sqrt{N}},$$

for h_k in a convex neighborhood of the true critical point $\widehat{h}_k^{\mathcal{R}}$, defined by $\nabla I_k(\widehat{h}_k^{\mathcal{R}}) = 0$. If we suppose that the estimated critical point $h_k^{\mathcal{R}}$, defined by $\nabla \Psi_k(h_k^{\mathcal{R}}) = 0$, lies in this neighborhood, and furthermore that the operator inequalities $0 < C_3 \leq \nabla^2 I_k \leq C_4$ hold on the same neighborhood, then one immediately derives the inequalities

$$\|h_k^{\mathcal{R}} - \widehat{h}_k^{\mathcal{R}}\| \leq \frac{C_2}{C_3 \sqrt{N}} \quad (32)$$

$$|\Psi_k(h_k^{\mathcal{R}}) - I_k(\widehat{h}_k^{\mathcal{R}})| \leq \frac{C_1}{\sqrt{N}} + \frac{C_2^2 C_4}{2C_3^2 N}, \quad (33)$$

which show convergence of $h_k^{\mathcal{R}}$ to $\widehat{h}_k^{\mathcal{R}}$ as $N \rightarrow \infty$.

The above discussion addresses the errors made at the k th time step of the algorithm. Further study is needed to understand how errors accumulate as k is iterated. The answer to this question will give guidance on how to distribute computational effort over the different time steps, and can be expected to parallel the same question as it arises for the LS algorithm.

To conclude this discussion, it is worthwhile to revisit the way in which our method of dynamic programming (finding H^B by induction over K steps backwards in time) leads to computational efficiency compared to a more direct approach which seeks to compute the optimal hedging strategy H^B simultaneously at all times. Fixing as before an R -dimensional subspace $\mathcal{R}(\mathcal{S})$ for the form of the hedging strategy at each time, direct optimization of a single convex function of $K \times R$ variables costs $\mathcal{O}(NR^2K^2)$ flops. By dynamic programming this is reduced to

K sequential optimizations of functions of R variables which will take $\mathcal{O}(NR^2K)$ flops. Accuracy is preserved by dynamic programming because the $KR \times KR$ Hessian matrix of the global optimization is approximately block diagonal over the individual time steps.

4 Numerical implementation

4.1 Geometric Brownian motion

We start with a one dimensional complete market in order to test the algorithm against well known exact solutions. Consider a market where the stock price process, discounted by the constant interest rate r , satisfies

$$\frac{dS_t}{S_t} = (\mu - r)dt + \sigma dW, \quad (34)$$

where μ and $\sigma > 0$ are constants and W is a one-dimensional P -Brownian motion. As is well known, the unique equivalent martingale measure Q has density dQ/dP given by the stochastic exponential of the constant market price of risk $\lambda = (\mu - r)/\sigma$ and the Merton portfolio for this market is given by

$$\hat{H}_t = \frac{\mu - r}{\gamma\sigma^2} \frac{1}{S_t}. \quad (35)$$

We can now compare the hedging portfolio “learned” by our algorithm with the “true” optimal hedging portfolio

$$H_t^B = \hat{H}_t + \mathcal{H}_t^B, \quad (36)$$

where \mathcal{H}_t^B is the Black–Scholes *delta hedging* portfolio replicating B . Similarly, the indifference prices calculated by the algorithm can be compared with the Black–Scholes price for the same claim.

We fix the parameters of the model at $S_0 = 1$, $\mu = 0.1$, $\sigma = 0.2$ and $r = 0.02$ over the period of one year $T = 1$. We apply the allocation algorithm with $N = 100000$ to two scenarios involving portfolio selection at discrete time intervals of $1/50$ (i.e. weekly): i) the Merton investment

problem; and ii) the hedging problem for the *writer* of a single at-the-money European put. For comparison to theory, we use the same Monte Carlo simulations, but rehedged weekly according to the theoretical formula (36). As for the subspace $\mathcal{R}(\mathcal{S})$, we use the three dimensional space spanned by the functions $\{1, s, s^2\}$.

Our results are displayed in figure 1, which shows the profit/loss distributions at time $T = 1$ for the true Merton (TM), true put option (TP), learned Merton (LM) and learned put option (LP) cases respectively. For comparison of their performances, Table 1 shows the mean and the standard deviation of these distributions in each of the four cases, as well as the final expected exponential utility with parameters $\gamma = 1/4, 1$ and 4 , corresponding to an increasing order of risk-aversion. As measures of the risk associated with each case, we also tabulate their value-at-risk and conditional value-at-risk for 90% (VaR_{90} and CVaR_{90}) and 99% (VaR_{99} and CVaR_{99}) confidence levels.

Case	Mean	St. Dev.	$\gamma = 1/4$	$\gamma = 1$	$\gamma = 4$	VaR_{99}	VaR_{90}	CVaR_{99}	CVaR_{90}
TM	0.1577	0.3963	-0.9661	-0.9238	-1.8561	-0.7629	-0.3504	-0.8940	-0.5367
LM	0.1596	0.3999	-0.9657	-0.9233	-1.8564	-0.7632	-0.3559	-0.8868	-0.5418
TP	0.0882	0.3957	-0.9830	-0.9902	-2.4569	-0.8330	-0.4206	-0.9653	-0.6072
LP	0.0901	0.3995	-0.9826	-0.9898	-2.4246	-0.8284	-0.4438	-0.9294	-0.6237

Table 1: Mean, standard deviation, final expected utilities and risk measures for the profit/loss distribution of the true Merton (TM), learned Merton (LM), true put option (TP), and learned put option (LP) portfolios for 100000 Monte Carlo simulations of stock prices following a geometric Brownian motion.

The learned estimate of the indifference price based on these 100000 simulations is 0.06944. Using the true (Black-Scholes delta hedging) strategy on the same simulated paths leads to the estimate 0.06935, while the theoretical Black-Scholes price is 0.06936. All the results for this example were obtained on a desktop PC in approximately 1 hour. While the estimates for the indifference prices are quite accurate in comparison to the Black-Scholes price, the hedge ratios on individual sample paths based on 100000 simulations are less satisfactory, due to large oscillations in time for the hedge coefficients calculated from the algorithm. In order to obtain more stable hedge coefficients (and consequently smoother hedge portfolios) without resorting to more elaborate smoothing procedures, we found necessary to increase the number

of simulations considerably. Figures 2 and 3 show the values of the hedge ratio along a single sample path calculated according to the true strategy and learned strategies for $N = 100000$ and $N = 1000000$. While up to $N = 100000$ the computational time increased linearly with the number of simulations, the run with $N = 1000000$ took approximately 24 hours to complete on the same desktop computer, likely due to larger memory requirements.

4.2 Stochastic Volatility

As an example of an incomplete market, consider a two factor stochastic volatility model of the form

$$\begin{aligned} dS_t &= S_t[(\mu - r)dt + \sqrt{Y_t}dW_t^1], \\ dY_t &= a(t, Y_t)dt + b(t, Y_t)[\rho dW_t^1 + \sqrt{1 - \rho^2}dW_t^2], \end{aligned} \quad (37)$$

where S_t is the price of a stock discounted by a constant interest rate r with instantaneous volatility $\sigma_t = \sqrt{Y_t}$. The processes S_t and Y_t are driven by the independent one-dimensional P -Brownian motions W_t^1 and W_t^2 and are assumed to have constant correlation $-1 < \rho < 1$. A popular choice for stochastic volatility models is to take Y_t to be a Cox–Ingersoll–Ross process [17]. In [13] we argue that exponential utility pricing in which the *reciprocal* of Y_t follows a CIR process leads to very efficient numerical computations of prices and hedge portfolios for the special case of claims which depend only on the terminal value of Y_t , such as a spot volatility put option. Of course our present method is not restricted to volatility claims, but in order to compare the results obtained by Monte Carlo simulations with those obtained in [13] we now adopt the specification of a reciprocal CIR process for Y_t , that is, we put

$$Y_t = \frac{(1 - \rho^2)(\mu - r)^2}{2R_t} \quad (38)$$

and take R_t evolving according to

$$dR_t = \alpha(\kappa - R_t)dt + \beta\sqrt{R_t} \left[\rho dW_t^1 + \sqrt{1 - \rho^2}dW_t^2 \right], \quad (39)$$

for constants $\alpha, \kappa, \beta > 0$ with $4\alpha\kappa > \beta^2$.

An application of the Itô formula then gives the following functional form for the coefficients in (37):

$$a(t, Y_t) = \alpha Y_t + \frac{2(\beta^2 - \alpha\kappa)}{(1 - \rho^2)(\mu - r)^2} Y_t^2, \quad (40)$$

$$b(t, Y_t) = - \left(\frac{2}{1 - \rho^2} \right)^{1/2} \frac{\beta}{(\mu - r)} Y_t^{3/2}. \quad (41)$$

We fix the CIR parameters to be $\alpha = 5, \beta = 0.04, \kappa = 0.001$, the stock parameters to be $S_0 = 1, \mu = 0.04, r = 0.02$, choose a constant correlation $\rho = 0.5$ and set the risk aversion parameter at $\gamma = 1$. With these values, the mean reversion time for Y_t is approximately two months and the long term distribution for the volatility $\sigma_t = \sqrt{Y_t}$ has expected value approximately 40%. To account for the portfolio dependence in both S_t and Y_t we took $\mathcal{R}(\mathcal{S})$ to be the six-dimensional space spanned by the functions $\{1, y, y^2, s, sy, s^2\}$.

We now apply the allocation algorithm to a volatility put option with payoff $(K - \sigma_T^2)^+$ and strike price $K = 0.15$. The indifference prices for such options can be quickly and accurately calculated as in [13] using a Fast Fourier Transform technique. We refer to these as the “exact” prices for purposes of comparison with the results obtained in the present paper. Due to the mean reversion in Y_t , options with maturity significantly longer than two months (our chosen mean reversion time) exhibit prices which are virtually independent on the current level of volatility. To obtain a more interesting behavior, we fix the time to maturity at $T = 0.2$ and compute the indifference prices with Y_0 varying in the interval $[0, 0.5]$ with equally spaced increments of size 0.01. Our results for $N = 1000$ and $N = 10000$ are displayed in figure 4. We observe that a new set of N Monte Carlo paths needs to be simulated for each initial value Y_0 . For $N = 1000$ the prices and hedging strategies corresponding to 50 different values of Y_0 were obtained in about 1 hour, whereas for $N = 10000$ computational time increased linearly to approximately 10 hours.

Next we consider a put option on the stock, that is, with payoff $(K - S_T)^+$. This is our first result where no exact solution for the indifference price is available for comparison. Instead, we try to reproduce the qualitative features observed in real markets for which stochastic volatility

regimes have been proposed. The table below shows the indifference prices obtained with $N = 10000$ for selected strike prices and times to maturity. Its last column contains the theoretical Black–Scholes price for the same model parameters and using the realized long term mean $E[\sigma_T] = E[\sqrt{Y_T}] = 0.4127$ (for $T = 0.8$) as a proxy for a constant volatility σ .

K	$T = 0.05$	$T = 0.1$	$T = 0.2$	$T = 0.4$	$T = 0.8$	BS price
0.8	0.0003	0.0025	0.0101	0.0265	0.0541	0.0506
0.9	0.0052	0.0140	0.0301	0.0548	0.0902	0.0882
1	0.0344	0.0487	0.0693	0.0981	0.1373	0.1375
1.1	0.1048	0.1141	0.1302	0.1567	0.1953	0.1972
1.2	0.1992	0.2012	0.2092	0.2288	0.2623	0.2659

Table 2: Indifference prices based on $N = 10000$ Monte Carlo simulations for a put option $(K - S_T)^+$ with different values of strike price and time to maturity, for a reciprocal affine stochastic volatility model.

Figure 5 shows the volatility surface obtained from the implied volatility for these indifference prices. The prices in the last two columns above are depicted in figure 6, which shows that the Black-Scholes prices for in-the-money put options are lower than the indifference prices, while the reverse occurs for out-of-the-money put options. The entire set of prices and hedging strategies for 5 different times to maturity and 17 different strike prices based on $N = 10000$ simulations was produced by our desktop PC in approximately 3.5 hours.

5 Discussion

This paper seeks to bridge the gap between the theory of exponential hedging in incomplete markets and the numerical implementation of that theory. Utility based hedging introduces several key concepts, notably certainty equivalent values and indifference prices which have no counterpart in complete markets. Therefore we have little experience or intuition on which to base our understanding of optimal trading in these markets. The simple and flexible Monte Carlo algorithm we introduce in this paper provides a test bed for realizing the theory of exponential hedging in essentially any market model. For example, problems involving American style early-

exercise options can in principle be easily included in our framework by following the Longstaff-Schwartz Monte Carlo method [22]. Using our method for a variety of problems should help one gain intuition and understanding of how exponential hedging works in practice and how it compares with other hedging approaches.

Our preliminary study of the geometric Brownian model shows not unexpectedly that the method performs better for pricing than hedging. Interestingly the indifference price, perhaps the key theoretical concept, appears to be better approximated than the two certainty equivalent values which define it. On the other hand, as we see from the sample path shown in Figure 2, the actual hedging strategy learned by the algorithm deviates a lot from the theoretical strategy along individual stock trajectories, and cannot be seen as reliable.

For the reciprocal affine stochastic volatility model, already at $N = 10000$ the algorithm produced accurate prices when compared with the explicit solution for pure volatility claims obtained in [13]. For the case of claims on the traded asset, for which no closed form solution is available, our algorithm generated option prices reproducing qualitative characteristics present in certain markets, such as implied volatility smiles or systematic biases when compared to the Black-Scholes prices for the same level of volatility.

Predictably, the basic method we use shows some distinct shortcomings which prevent it from being taken as a *de jure* guide to real trading. Approximation one, arising by restricting possible hedge strategies to a low dimensional subspace, clearly will often lead to unsuitable strategies. An improvement in this respect would be to use the known leading order terms for a given problem, say the Black-Scholes delta hedging strategy, as an input in the algorithm, which would then search for higher order ‘corrections’ within the subspace \mathcal{R} . Another difficulty we noticed arising in our method is that learned strategies fluctuate far too much in time. Some simple smoothing procedure in time might lead to a marked improvement in hedging. Concerning approximation two, the finite sample size error, a brief study of the size of the constants which enter the (pessimistic) estimates (32) and (33) suggests that reliable learned strategies will demand a very large value of N (in our simulations, $N = 10000$ gave reliable prices, but not hedging strategies). The possibility of using variance reduction techniques (see e.g. [12]) should

be considered an area for further study.

Putting aside the obvious drawbacks of the algorithm, we can see that our very simple and direct method will shed light on most conceptual difficulties arising in exponential hedging in incomplete markets. It implements the spirit of dynamic programming and prices claims quite reliably, even if it cannot easily produce accurate estimates of hedging strategies. On these merits alone, we think our algorithm deserves much further study and refinement.

References

- [1] J. Abad and T. R. Hurd. *Error bounds for Monte Carlo based portfolio optimization*. Working paper available at www.math.mcmaster.ca/tom, 2003.
- [2] D. Becherer. *Rational Hedging and Valuation with Utility-based Preferences*. PhD thesis, Technischen Universität Berlin, 2001.
- [3] D. Becherer. Utility-indifference hedging and valuation via reaction-diffusion systems. *Proc. Roy. Soc.* 460, 2041, 27–51, 2004.
- [4] P. Boyle, M. Broadie, and P. Glasserman. Monte Carlo methods for security pricing. In *Option pricing, interest rates and risk management*, Handb. Math. Finance, pages 185–238. Cambridge Univ. Press, Cambridge, 2001.
- [5] E. Clément, D. Lamberton, and P. Protter. An analysis of a least squares regression method for American option pricing. *Finance Stoch.*, 6(4):449–471, 2002.
- [6] J. Cvitanic, L. Goukasian, and F. Zapatero. Monte Carlo computation of optimal portfolios in complete markets. Preprint, 2000.
- [7] J. Cvitanić, W. Schachermayer, and H. Wang. Utility maximization in incomplete markets with random endowment. *Finance Stoch.*, 5(2):259–272, 2001.

- [8] M. H. A. Davis. Option pricing in incomplete markets. In *Mathematics of derivative securities (Cambridge, 1995)*, volume 15 of *Publ. Newton Inst.*, pages 216–226. Cambridge Univ. Press, Cambridge, 1997.
- [9] F. Delbaen, P. Grandits, T. Rheinländer, D. Samperi, M. Schweizer, and C. Stricker. Exponential hedging and entropic penalties. *Math. Finance*, 12(2):99–123, 2002.
- [10] F. Delbaen and W. Schachermayer. A general version of the fundamental theorem of asset pricing. *Math. Ann.*, 300(3):463–520, 1994.
- [11] J. Detemple, R. Garcia, and M. Rindisbacher. A Monte Carlo method for optimal portfolios. To appear in *Journal of Finance*, 2003.
- [12] P. Glassermann. *Monte Carlo Methods in Financial Engineering* Springer–Verlag, 2003.
- [13] M. R. Grasselli and T. R. Hurd. Indifference pricing and hedging in stochastic volatility models. Submitted to *Math. Finance*, 2004.
- [14] J. M. Harrison and S. R. Pliska. Martingales and stochastic integrals in the theory of continuous trading. *Stochastic Process. Appl.*, 11(3):215–260, 1981.
- [15] V. Henderson. Valuation of Claims on Non-Traded Assets using Utility Maximization. *Mathematical Finance*, 12:4, 351–373, 2002.
- [16] V. Henderson. Valuing Real Options Without a Perfect Spanning Asset. Preprint, 2004.
- [17] S. L. Heston. A closed–form solution for options with stochastic volatility with applications to bond and currency options. *Rev. Financial Studies* 6, 2, 327–343, 1993.
- [18] S. D. Hodges and A. Neuberger. Optimal replication of contingent claims under transaction costs. *Rev. Fut. Markets*, 8:222–239, 1989.
- [19] I. Karatzas, J. P. Lehoczky, S. E. Shreve, and G.-L. Xu. Martingale and duality methods for utility maximization in an incomplete market. *SIAM J. Control Optim.*, 29(3):702–730, 1991.

- [20] I. Karatzas and S. E. Shreve. *Methods of mathematical finance*, volume 39 of *Applications of Mathematics*. Springer-Verlag, New York, 1998.
- [21] D. Kramkov and W. Schachermayer. The asymptotic elasticity of utility functions and optimal investment in incomplete markets. *Ann. Appl. Probab.*, 9(3):904–950, 1999.
- [22] F. A. Longstaff and E. S. Schwartz. Valuing American options by simulation: a simple least-square approach. *Rev. Fin. Studies*, 14:113–147, 2001.
- [23] C. Munk. The valuation of contingent claims under portfolio constraints: reservation buying and selling prices. *Euro. Fin. Rev.*, 3:347-388, 1999.
- [24] M. P. Owen. Utility based optimal hedging in incomplete markets. *Ann. Appl. Probab.*, 12(2):691–709, 2002.
- [25] P. Protter. *Stochastic integration and differential equations*, volume 21 of *Applications of Mathematics*. Springer-Verlag, Berlin, 1990. A new approach.
- [26] W. Schachermayer. Optimal investment in incomplete markets when wealth may become negative. *Ann. Appl. Probab.*, 11(3):694–734, 2001.
- [27] V. Young and T. Zariphopoulou. Pricing dynamic insurance risks using the principle of equivalent utility. *Scand. Actuarial J.*, 4:246–279, 2002.
- [28] T. Zariphopoulou. A solution approach to valuation with unhedgeable risks. *Finance Stoch.* 5, 1, 61–82, 2001.




Figure 1: Profit/loss distributions at time $T = 1$ year for Merton's portfolio and the optimal hedge portfolio for the writer of a put option using both the theoretical and learned trading strategies based 100000 simulations of a stock following geometric Brownian motion.




Figure 2: The hedge amounts (number of shares held) for the writer of one put option on a simulated sample path (dotted line) of duration one year of a stock following geometric Brownian motion. The solid line shows the strategy learned with $N=100000$; the broken line shows the theoretical Black-Scholes-Merton strategy.




Figure 3: The hedge amounts (number of shares held) for the writer of one put option on a simulated sample path (dotted line) of duration one year of a stock following geometric Brownian motion. The solid line shows the strategy learned with $N=1000000$; the broken line shows the theoretical Black-Scholes-Merton strategy.




Figure 4: Indifference prices for volatility put options $(K - \sigma_T^2)^+$ with strike price $K = 0.15$ and maturity $T = 0.2$ for different initial values Y_0 , calculated using 1000 and 10000 Monte Carlo simulations of a reciprocal CIR process for Y_t . The blue line correspond to the same prices calculated from the solution of the HJB equation associated with this model.




Figure 5: Implied volatility surface for put options $(K - S_T)^+$ obtained from the indifference prices based on 10000 Monte Carlo simulations of a reciprocal CIR stochastic volatility model.




Figure 6: Indifference price for the writer of put options $(K - S_T)^+$ with $T = 0.8$ based on 10000 Monte Carlo simulations of a reciprocal CIR stochastic volatility model compared with the Black-Scholes price calculated using the realized mean $E[\sigma_T] = 0.4127$ as a proxy for a constant volatility.