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

First version: November 25, 2002
Current version: November 25, 2002

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 family of exponential utility functions $U(x) = -\frac{e^{-\gamma x}}{\gamma}$, $\gamma > 0$. 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 much 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 portfolio theory provides a coherent, general and economically sound approach to risk–management in general financial models. This theory is built on the principle that market agents invest rationally by seeking to maximize their expected utility over some time period, where their utility function encodes the "happiness" they derive in holding a given level of wealth. Key works in this program are those of [18], [19], [20]. The culmination of these results is a body of theory which give necessary and sufficient conditions for existence and uniqueness of optimal portfolios in a broad range of contexts.

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. 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. This framework leads to new concepts, notably the Davis price [9] and the indifference price of the contingent claim [16].

This much more general theory is naturally applicable in areas such as insurance where the complete market theory appears inappropriate [28]. In this 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.

Practical implemention 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 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 tree and lattice based methods. A third broad approach can be called Monte Carlo or random simulation based 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 [3], have not been extensively used for optimal portfolio theory. Some works related to this in the context of complete markets are [13] and [7]. 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 [21]. 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 quite difficult to learn from the abstract theory.

The paper is organized as follows. Section 2 provides the reader with a rather detailed survey of the current theory of optimal portfolios. We give careful statements of the main results concerning the existence and uniqueness of optimal solutions for Merton's problem. We also review the framework of utility based hedging, introducing the key concepts and the basic existence/uniqueness results. The special case of exponential utility is discussed in some detail, because it has the important property that optimal solutions are independent of the level of wealth. This property has an important implication for our proposed Monte Carlo algorithm.

Section 3 focuses on the dynamics of portfolio optimization, in particular, the principle of dynamic programming. The concepts of certainty equivalent value, indifference price and the Davis price are introduced. The example of the geometric Brownian motion market is worked out in some detail. Section 4 specializes to the discrete time hedging framework and gives explicit formulas for dynamic programming.

The main innovation of the paper is the exponential utility algorithm given in section 5. It is a Monte Carlo method for learning the optimal trading strategy for the class of discrete time hedging problems introduced in section 4. This algorithm is inspired by the least-squares algorithm of Longstaff and Schwartz for pricing American options. Interestingly, our method works well only for the expopnenial utility, and no simple extension suggests itself for general utility functions. Section 6 describes our first application of the algorithm to hedging in a one-dimensional geometric Brownian motion model. We focus on this exactly solvable model in order to have explicit formulas with which to compare our Monte Carlo simulation. While the hedging strategies learned by the algorithm are somewhat crude, we find that the computed indifference prices are quite accurate. In our concluding section 7, we discuss the various advantages and drawbacks we see in the method.

2 Utility based hedging for semimartingale markets

The hedging problem is the problem of a market agent who faces a liability B at a time T and must invest in the market over the period $[0,T]$ in an efficient, rational or otherwise optimal way to reduce the risk of the liability. The randomness of the market is represented by a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\in[0,T]}, P)$ satisfying the "usual conditions" of right continuity and completeness and we assume for simplicity that $\mathcal{F} = \mathcal{F}_T$. The discounted prices of tradeable assets in the market are given by the $\mathbb{R}^d$–valued *càdlàg* semimartingale $S_t = (S^1_t, \dots, S^d_t)$ on the filtration $(\mathcal{F}_t)$. The liability B is assumed to be an $\mathcal{F}_T$–measurable random variable.

A portfolio process, or a trading strategy, is an $\mathbb{R}^d$-valued predictable S–integrable process $H_t = (H^1_t, \dots, H^d_t)$, which represents the agent's asset allocations, that is, how many units of each traded asset are held by the agent at each time t. The class of such processes is denoted by $L(S)$ [23]. We assume that the portfolio is *self-financing* (i.e. the changes in its discounted market value are solely due to the random changes in the prices of the traded assets) so that the agent's discounted wealth at each time t is given by the process

$$X_t = x + (H \cdot S)_t := x + \int_0^t H_u dS_u, \qquad t \in [0,T],$$

where $x \in \mathbb{R}$ is some deterministic initial wealth.

To rule out strategies for which the wealth assumes arbitrarily negative values (such as "doubling strategies"), we need to assume some admissibility condition on the possible portfolio processes. Following [15], we say that

Definition 2.1 *The class $\mathcal{H}$ of admissible portfolios consists of the process $H \in L(S)$ for which $(H \cdot S)_t$ is P–a.s. uniformly bounded from below.*

More explicitly, H is admissible if there exists a constant $k \geq 0$ (possibly depending on H, but neither on t nor on ω) such that

$$(H \cdot S)_t(\omega) \geq -k,$$

for almost all $\omega \in \Omega$ and all $t \in [0,T]$.

As a first consequence of this notion of admissibility, we have the following useful result concerning the closedness of the class of local martingales under stochastic integration [11, theorem 2.9]:

Lemma 2.2 *If S is a local martingale and H is an admissible integrand for S, then $(H \cdot S)$ is a local martingale. Consequently, $(H \cdot S)$ is a supermartingale.*

Regarding martingale measures, we adopt the following definition.

Definition 2.3 *A probability measure Q is called an absolutely continuous (resp. equivalent) local martingale measure for S if $Q \ll P$ (resp. $Q \sim P$) and S is a local martingale under Q.*

We denote the set of absolutely continuous (resp. equivalent) local martingale measures for the price process S by $\mathcal{M}^a(S)$ (resp. by $\mathcal{M}^e(S)$). Observe that, due to lemma 2.2, a probability measure $Q \ll P$ (resp. $Q \sim P$) is an absolutely continuous (resp. equivalent) local martingale measure if and only if $(H \cdot S)$ is a local martingale under Q for any $H \in \mathcal{H}$.

To ensure a viable market, free of arbitrage, we assume the technical condition "No Free Lunch with Vanishing Risk" (NFLVR), which is slightly more general than "No Arbitrage" (NA). The reader is referred to [11, sections 2 and 3] for the precise definition of these notions, as well as the relations between them. In its most general form [12], the fundamental theorem of asset pricing (FTAP) asserts the equivalence between (NFLVR) and the existence of an equivalent σ–martingale measure for the price process S [12], which might fail to be in $\mathcal{M}^e(S)$ if we allow S to have unbounded unpredictable jumps. The technicality of using σ–martingales can be avoided, however, if we restrict ourselves to price processes S which are *locally bounded*. By that we mean that there exists a localizing sequence of stopping times $\{T_n\}$ such that, for each n, the stopped processes S^{T_n} are bounded. In this context, we have [11, corollary 1.2]:

Theorem 2.4 (FTAP) *If S is a locally bounded semimartingale, then there exists an equivalent local martingale measure Q for S if and only if S satisfies (NFLVR).*

In view of this theorem, we will henceforth assume that S is locally bounded and that

Assumption 1 (NFLVR) $\mathcal{M}^e(S) \neq \emptyset$.

The hedging problem can be made specific by introducing the agent's utility $U : \mathbb{R} \to \mathbb{R} \cup \{-\infty\}$, a concave, strictly increasing, differentiable function. Beginning with initial capital $x \in \mathbb{R}$, the agent then solves the optimal hedging problem

$$\sup_{H \in \mathcal{H}} E\left[U\left(x + (H \cdot S)_T - B\right)\right]. \tag{1}$$

If $B \equiv 0$, the optimal hedging problem reduces to Merton's optimal investment problem

$$\sup_{H \in \mathcal{H}} E\left[U\left(x + (H \cdot S)_T\right)\right]. \tag{2}$$

To assert the existence and uniqueness of solutions to problems of the form (2) in incomplete markets, one first needs to impose further technical restrictions on the class of utility functions. In the next assumption we summarize the main properties required to hold throughout this paper. They include the "reasonable asymptotic elasticity" condition as defined in [26, definition 1.5].

Assumption 2 *The utility function $U : \mathbb{R} \to \mathbb{R} \cup \{-\infty\}$ is increasing on $\mathbb{R}$, continuous on $\{U > -\infty\}$, differentiable and strictly concave on dom$(U) =$ int$\{U > -\infty\}$, satisfying*

$$\lim_{x\to\infty} U'(x) = 0. \tag{3}$$

Furthermore, we assume that one of the following cases hold.

Case 1: *dom$(U) = (0,\infty)$, with $\lim_{x\to 0} U'(x) = \infty$ and $\limsup_{x\to\infty} \frac{xU'(x)}{U(x)} < 1$.*

Case 2: *dom$(U) = \mathbb{R}$, with $\lim_{x\to -\infty} U'(x) = \infty$,*

$$\limsup_{x\to\infty} \frac{xU'(x)}{U(x)} < 1 \quad \textit{and} \quad \liminf_{x\to-\infty} \frac{xU'(x)}{U(x)} > 1.$$

The central technical weaponry used to address the general solution to problem (2) is convex duality, by means of which the utility maximization problem over 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"). The first step is to define the conjugate function V as the Legendre transform of the function $-U(-x)$, that is

$$V(y) := \sup_{x\in\mathbb{R}} [U(x) - xy], \qquad y > 0. \tag{4}$$

It follows from well known results in convex analysis [24], that the function V has the properties listed below.

Proposition 2.5 *If U satisfies assumption 2, then the conjugate function V is finite valued, differentiable, strictly convex on $(0,\infty)$ and satisfies*

$$\lim_{y\to 0} V(y) = \lim_{x\to\infty} U(x), \qquad \lim_{y\to 0} V'(y) = -\infty. \tag{5}$$

Moreover, the behaviour of V at infinity is determined by the two cases in assumption 2 as follows:

Case 1: $\lim_{y\to\infty} V(y) = \lim_{x\to 0} U(x)$ *and* $\lim_{y\to\infty} V'(y) = 0$.

Case 2: $\lim_{y\to\infty} V(y) = \infty$ *and* $\lim_{y\to\infty} V'(y) = \infty$.

Both the primal and dual problems are solved over different domains depending on which of the two cases above we are dealing with. We start with the first case, for which the present state–of–the–art solution can be found in [20]. Since in this case the utility function is only defined for positive wealths, we need to consider the set

$$\mathcal{X}(x) = \{X \geq 0 : X_t = x + (H \cdot S)_t, \text{ for some } H \in L(S), 0 \leq t \leq T\}. \quad (6)$$

It is clear that $x+(H\cdot S)_t \geq 0$ implies that the portfolio H must be admissible, that is

$$\mathcal{X}(x) \subset \{X_t = x + (H \cdot S)_t, H \in \mathcal{H}, t \in [0, T]\}$$

with a strict inclusion. Next we move from the set of processes $\mathcal{X}(x)$ to the set of positive random variables

$$C(x) = \{g \in L^0_+(\Omega, \mathcal{F}_T, P) : g \leq X_T, \text{ for some } X \in \mathcal{X}(x)\} \quad (7)$$

and observe that, since the utility function is increasing, the primal problem for case 1 written in the form

$$\sup_{X \in \mathcal{X}(x)} E\left[U(X_T)\right]. \quad (8)$$

is equivalent to

$$u(x) = \sup_{g \in C(x)} E\left[U(g)\right]. \quad (9)$$

At this point, in order to exclude trivial cases, we make the following assumption.

Assumption 3 *The value function u defined in (9) satisfies $u(x) < \infty$, for some $x > 0$.*

As for the domain of the dual problem, one looks for a set with the property of being in a "polar relation" with the set C (the reader is refered to [4] for the definition of the polar of a subset of $L^0_+(\Omega, \mathcal{F}, P)$). In the mathematical finance literature [11, 18], the sets $\mathcal{M}^e(S)$ and $\mathcal{M}^a(S)$ were considered. One of the main technical novelties in [20] was to enlarge this domain in order to obtain a set D in "perfect" polar relation with C (see [20, proposition 3.1]). The set D turns out to be the convex, solid hull of $\mathcal{M}^a(S)$ in $L^0_+(\Omega, \mathcal{F}, P)$ (topologized by convergence in measure). Amongst the several equivalent characterizations of the set D, we single out the following [25]

$$\begin{aligned} D \;\; = \;\; & \Big\{Y_T \in L^0_+(\Omega, \mathcal{F}_T, P) : \text{there exists a sequence} \\ & (Q_n)_{n=1}^{\infty} \in \mathcal{M}^a(S) \text{ such that } Y_T \leq (a.s.) \lim_{n\to\infty} \frac{dQ_n}{dP}\Big\} \end{aligned} \quad (10)$$

For $y > 0$, let us define $D(y) = yD$. The dual problem for case 1 can now be formulated as

$$v(y) = \inf_{Y_T \in D(y)} E[V(Y_T)]. \tag{11}$$

The next theorem states the existence and uniqueness of solution for the problems (9) and (11) for utilities restricted to positive wealth [20, theorem 2.2].

Theorem 2.6 *Suppose that assumptions 1, 2 (case 1) and 3 are satisfied. Then, for any $x \in dom(U)$ and $y > 0$, the problems*

$$u(x) = \sup_{X_T \in C(x)} E[U(X_T)], \qquad v(y) = \inf_{Y_T \in D(y)} E[V(Y_T)] \tag{12}$$

have unique optimizers $\widehat{X}_T(x) \in C(x)$ and $\widehat{Y}_T(y) \in D(y)$ satisfying

$$U'(\widehat{X}_T(x)) = \widehat{Y}_T(y), \tag{13}$$

where x and y are related by $u'(x) = y$.

We note that this theorem and its proof apply unchanged if we modify case 1 of assumption 2 to allow for utilities defined on an interval of the form (a, ∞), for any $a \in \mathbb{R}$, provided we impose that $\lim_{x \to a} U'(x) = \infty$. Observe also that the optimizer $\widehat{X}_T(x)$ can be uniquely expressed as

$$\widehat{X}_T(x) = x + (\widehat{H}(x) \cdot S)_T,$$

for $\widehat{H}(x) \in \mathcal{H}$, whereas the optimizer $\widehat{Y}_T(y)$, even for the cases where $\widehat{Y}_T(y)/y$ is *not* the density of an absolutely continuous martingale measure (by having its total P–mass strictly less than 1), can be arbitrarily approximated by elements in $\mathcal{M}^a(S)$ (in the sense of almost sure convergence).

For utility functions defined on the entire real line the problem is more involved, due to the fact that the class of admissible portfolios as in definition 2.2 turns out to be too narrow to contain the optimal solution. One approach is to start with the dual problem, for which [2] shows that an optimal solution always exists (under very general conditions). Then the set of allowed portfolios can be characterized in terms of it. This opens up a plethora of definitions of "allowed" portfolios. The reader interested in this line of thought is referred to [10, 17], where the exponential utility is addressed, and to [27] for more general utility functions.

A more direct idea is to concentrate on random variables which do not necessarily arise as terminal values of wealth processes for any portfolios, but which can be arbitrarily approximated by such objects. Different such domains of optimization over random variables have been proposed [26, 14], the difference being the kind of topology (convergence) adopted to describe the approximation mentioned above. In what follows, we adopt the approach proposed in [26], and specialize later on to the case of exponential utility $U(x) = -\frac{e^{-\gamma x}}{\gamma}$, where sharper results can be quoted.

We denote by $C_U^b(x)$ the class of random variables which have integrable utility and can be dominated by the terminal wealth of admissible portfolios, that is,

$$\begin{aligned} C_U^b(x) &= \{g \in L^0(\Omega, \mathcal{F}_T, P) : g \leq x + (H \cdot S)_T \\ &\quad \text{for some } H \in \mathcal{H} \text{ and } U(g) \in L^1(\Omega, \mathcal{F}_T, P)\}. \end{aligned} \tag{14}$$

Next consider the closure of the set $\{U(g) : g \in C_U^b(x)\}$ in the topology of $L^1(\Omega, \mathcal{F}_T, P)$. Putting $U(\infty) := \lim_{x\to\infty} U(x)$, we see that the the utility function is a bijection between $\mathbb{R}$ and $\mathbb{R}$, if $U(\infty) = \infty$, and a bijection between $\mathbb{R} \cup \{\infty\}$ and $(-\infty, U(\infty)]$ otherwise. We can therefore write a general element in this closure as $U(f)$ for some $f \in L^0(\Omega, \mathcal{F}_T, P; \mathbb{R} \cup \{\infty\})$. The set of such random variables is denoted by $C_U(x)$, that is,

$$\begin{aligned} C_U(x) &= \big\{f \in L^0(\Omega, \mathcal{F}_T, P; \mathbb{R} \cup \{\infty\}) : U(f) \text{ is in the} \\ &\quad L^1(P)\text{-closure of } \{U(g) : g \in C_U^b(x)\}\big\}. \end{aligned} \tag{15}$$

The primal optimization problem then becomes

$$u(x) = \sup_{f \in C_U(x)} E[U(f)]. \tag{16}$$

As in case 1, to exclude trivial cases we make the following assumption.

Assumption 4 *The value function u defined in (16) satisfies $u(x) < U(\infty)$, for some $x \in \mathbb{R}$.*

Complicated as the domain $C_U(x)$ might seem, the good news is that in this setting the optimization domain for the dual problem is simply $\mathcal{M}^a(S)$, as opposed to the enlarged set D of case 1. That is, the dual problem is now

$$v(y) = \inf_{Q \in \mathcal{M}^a(S)} E\left[V\left(y\frac{dQ}{dP}\right)\right]. \tag{17}$$

We can now state a theorem for case 2 [26, theorem 2.2].

Theorem 2.7 *Suppose that assumptions 1, 2 (case 2) and 4 are satisfied. Then:*

1. *For any $x \in \mathbb{R}$ and $y > 0$, the problems*

$$u(x) = \sup_{f \in C_U(x)} E[U(f)], \qquad v(y) = \inf_{Q \in M^a(S)} E\left[V\left(y\frac{dQ}{dP}\right)\right] \tag{18}$$

have unique optimizers $\widehat{f}(x) \in C_U(x)$ and $\widehat{Q}(y) \in \mathcal{M}^a(S)$ satisfying

$$U'(\widehat{f}(x)) = y\frac{d\widehat{Q}(y)}{dP}, \tag{19}$$

where x and y are related by $u'(x) = y$.

2. *If it occurs that $\widehat{Q}(y) \in \mathcal{M}^e(S)$, then $\widehat{f}(x)$ equals the terminal value $\widehat{X}_T(x)$ of a uniformly integrable $\widehat{Q}(y)$-martingale of the form*

$$\widehat{X}_t(x) = x + (\widehat{H}(x) \cdot S)_t, \quad t \in [0, T],$$

for some $\widehat{H}(x) \in L(S)$.

Observe that the optimizer $\widehat{f}(x) \in C_U(x)$ does not need to be the terminal wealth of any portfolio. However, by construction of the set $C_U(x)$, its utility can be arbitrarily approximated by the utility of terminal wealth of admissible portfolios. As for the optimizer of the dual problem, recall from proposition 2.5 that $\lim_{y \to 0} V(y) = U(\infty)$. Therefore for all cases when $U(\infty) = \infty$, the minimizer must satisfy $\frac{d\widehat{Q}(y)}{dP} > 0$ almost surely, implying that $\widehat{Q}(y) \in \mathcal{M}^e(S)$ and item 2 holds. In such cases, $\widehat{f}(x)$ itself can be achieved by trading according to a portfolio $\widehat{H} \in L(S)$. Although $\widehat{H}$ might not be in $\mathcal{H}$, the wealth process generated by it, being a uniformly integrable $\widehat{Q}(y)$ martingale, certainly does not arise from a "doubling strategy", so that $\widehat{H}$ can be considered *a posteriori* to be an "allowed" portfolio. Turning this argument around was the starting point of the aforementioned approaches to extend the domain of the primal problem to include such portfolios [10, 17, 27]

But the minimizer $\widehat{Q}(y)$ is also equivalent to P in other cases, and it is here that we specialize to an exponential utility of the form $U(x) = -\frac{e^{-\gamma x}}{\gamma}$, $\gamma > 0$. Observe that for this utility we have

$$\limsup_{x \to \infty} \frac{xU'(x)}{U(x)} = -\infty < 1$$

and

$$\liminf_{x \to -\infty} \frac{xU'(x)}{U(x)} = \infty > 1,$$

so that it satisfies all the conditions for case 2 of assumption 2. Observe further that its dual function is

$$V(y) = \frac{y}{\gamma}(\log y - 1),$$

so that the dual problem (17) is equivalent to the problem of finding a measure in $\mathcal{M}^a(S)$ with minimal relative entropy with respect to P, that is,

$$\inf_{Q \in \mathcal{M}^a(S)} E\left[\frac{dQ}{dP}\log\left(\frac{dQ}{dP}\right)\right]. \tag{20}$$

It follows from [6] that the minimizer of this problem (which incidentally is independent of y) will be an equivalent local martingale measure provided there exists at least one measure in $\mathcal{M}^e(S)$ with finite relative entropy, allowing us to use item 2 of theorem 2.7.

Corollary 2.8 *Let $U(x) = -\frac{e^{-\gamma x}}{\gamma}$, $\gamma > 0$, and suppose that assumptions 1 and 4 hold. If in addition we have that*

$$E\left[\frac{dQ}{dP}\log\left(\frac{dQ}{dP}\right)\right] < \infty, \tag{21}$$

for some $Q \in \mathcal{M}^e(S)$, then the minimizer $\widehat{Q}(y)$ of theorem 2.7 is the equivalent local martingale measure $\widehat{Q}$, independent of $y > 0$, which minimizes the relative entropy with respect to P among all absolutely continuous martingale measures. Therefore $\widehat{f}(x)$ equals the terminal value $\widehat{X}_T(x)$ of a uniformly integrable $\widehat{Q}$-martingale of the form

$$\widehat{X}_t(x) = x + (\widehat{H}(x)\cdot S)_t,$$

for some $\widehat{H}(x) \in L(S)$.

We now move to the subject of solving the hedging problem (1). Once more the solutions will take place in different domains and involve different techniques depending on whether our utility function falls into case 1 or case 2 of assumption 2. In either case, we are going to assume that the random claim that we want to hedge is a bounded random variable.

Assumption 5 $B \in L^\infty(\Omega, \mathcal{F}_T, P)$.

We start with the first case, which was solved in [8]. Observe that to account for the presence of a random claim at time T, it is not enough to consider *positive* random variables which are dominated by terminal values of admissible portfolios, as was done in (7). We therefore consider the set

$$\mathcal{C}(x) = \{g \in L^0(\Omega, \mathcal{F}_T, P) : g \leq x + (H \cdot S)_T, \text{ for some } H \in \mathcal{H}\}. \tag{22}$$

The primal problem now becomes

$$u(x) = \sup_{g \in \mathcal{C}(x)} E[U(g - B)], \tag{23}$$

where it is understood that $U(x) = -\infty$ whenever $x \leq 0$.

As in the previous cases, we assume the following.

Assumption 6 *The value function u defined in (23) satisfies $|u(x)| < \infty$ for some $x > \|B\|_\infty$.*

Recall that the crucial point in the proof of theorem 2.6 was the use of the polar relation between the sets C and D as subsets of $L^0_+(\Omega, \mathcal{F}_T, P)$, for which a version of the bipolar theorem can be used [4]. In the absence of such results for subsets of $L^0(\Omega, \mathcal{F}_T, P)$ as a whole, we are led to consider an appropriate subset of $L^\infty(P)$, namely

$$\mathcal{C} = \mathcal{C}(0) \cap L^\infty(\Omega, \mathcal{F}_T, P). \tag{24}$$

Accordingly, to obtain a perfect polar relation, we need to modify the definition for the domain of the dual problem. The natural space to define the polar of a subset of L^∞ is its topological dual $(L^\infty)^*$. We therefore define

$$\mathcal{D} = \{Q \in (L^\infty)^* : \|Q\| = 1 \text{ and } Q(g) \leq 0 \text{ for all } g \in \mathcal{C}\}. \tag{25}$$

To obtain a more concrete characterization of this set, notice that $\mathcal{D} \in (L^\infty)^*_+$ (since $\mathcal{C}$ contains all the *negative* bounded random variables). The good news about the set $(L^\infty)^*_+$ is that it can be identified with the set af all nonnegative *finitely* additive bounded set functions on $\mathcal{F}_T$ which vanish on the P–null sets. Moreover, any such function $Q \in (L^\infty)^*_+$ can be uniquely decomposed into its regular part Q^r and its singular part Q^s as follows

$$Q = Q^r + Q^s,$$

where $Q^r \geq 0$ is *countably* additive and $Q^s \geq 0$ is *purely finitely* additive. Naturally, Q^r corresponds to a measure which is absolutely continuous with respect to P and whose Radon–Nikodym derivative is denoted by $\frac{dQ^r}{dP}$. We now look at the subset of regular elements in $\mathcal{D}$, namely

$$\mathcal{D}^r = \{Q \in \mathcal{D} : Q^s = 0\} = \mathcal{D} \cap L^1(\Omega, \mathcal{F}_T, P). \tag{26}$$

Since all elements in $\mathcal{D}$ have unit norm, it follows that $\mathcal{D}^r$ consists of *probability* measures which are absolutely continuous with respect to P. In fact, since we are assuming that the processes S are locally bounded, it can be shown that $\mathcal{D}^r$ is nothing but our familiar $\mathcal{M}^a(S)$, the set of absolutely continuous local martingale measures for S [2, lemma 1.1 (b)]. The dual problem in this case is

$$v(y) = \inf_{Q\in\mathcal{D}} \left\{ E\left[V\left(y\frac{dQ^r}{dP}\right) - y\frac{dQ^r}{dP}B\right] - yQ^s(B)\right\}, \tag{27}$$

where one should notice that the domain of optimization is the entire $\mathcal{D}$, with the dual function V contributing to it only through its regular subset $\mathcal{D}^r$, whereas the dependence on the claim B is manifested on both its regular and singular parts. In this respect, it is worth mentioning that our old set D (for Merton's problem) can also be characterized as the regular part of the weak-star closure of the convex solid hull of $\mathcal{M}^a(S)$ in $(L^\infty)^*$ (whose elements can have total P–mass strictly less than 1). From this perspective, it becomes clear that the extra care necessary to treat the hedging problem in this case comes from dealing with both the regular and singular parts of elements in the domain of the dual problem. The main result in this case is [8, theorem 3.1]

Theorem 2.9 *Suppose that assumptions 1, 2 (case 1), 5 and 6 are satisfied. Let $x_0 = \sup_{Q\in\mathcal{D}} Q(B)$. Then, for any $y > 0$, the dual problem*

$$v(y) = \inf_{Q\in\mathcal{D}} \left\{ E\left[V\left(y\frac{dQ^r}{dP}\right) - y\frac{dQ^r}{dP}B\right] - yQ^s(B)\right\} \tag{28}$$

has a unique (up to singular part) optimizer $\widehat{Q}(y) \in \mathcal{D}$ and, for any $x > x_0$, the primal problem

$$u(x) = \sup_{X_T\in\mathcal{C}(x)} E[U(X_T - B)] \tag{29}$$

has unique optimizer $\widehat{X}_T(x) \in \mathcal{C}(x)$ satisfying

$$U'(\widehat{X}_T(x) - B) = y\frac{d\widehat{Q}^r(y)}{dP}, \tag{30}$$

where x and y are related by $u'(x) = y$.

Regarding the second case of assumption 2, the optimal hedging problem has been solved in [10] for the exponential utility and claims B satisfying a boundedness conditions weaker than assumption 5. In [22], the problem was solved for general utility functions with reasonable asymptotic elasticity (which include the exponential) but bounded claims (although some remarks are offered on how to extend the result to possibly unbounded ones). We describe here the solution of [22], since it follows the same techniques of [26] and [8], for which we have already developed most of the notation. In the presence of a claim satisfying assumption 5, the analogue of the set C_U^b defined in (14) is

$$\begin{aligned} \mathcal{C}_U^b(x) &= \{g \in L^0(\Omega, \mathcal{F}_T, P) : g \leq x + (H \cdot S)_T - B \\ &\quad \text{for some } H \in \mathcal{H} \text{ and } U(g) \in L^1(\Omega, \mathcal{F}_T, P)\}. \end{aligned} \tag{31}$$

Similarly, we replace the set $C_U(x)$ by

$$\begin{aligned} \mathcal{C}_U(x) &= \big\{f \in L^0(\Omega, \mathcal{F}_T, P; \mathbb{R} \cup \{\infty\}) : U(f - B) \text{ is in the} \\ &\quad L^1(P)\text{-closure of } \{U(g) : g \in C_U^b(x)\}\big\}. \end{aligned} \tag{32}$$

The interpretation of this set is the same as before, only this time we have to account for the random claim B. Namely, it consists of random variables which, after subtracting the claim B, have a utility that can be arbitrarily approximated by the utility of terminal wealth of admissible portfolios less the claim B.

Our modified primal problem now reads

$$u(x) = \sup_{f \in \mathcal{C}_U(x)} E[U(f - B)], \tag{33}$$

for which we assume the following.

Assumption 7 *The value function u defined in (33) satisfies $u(x) < U(\infty)$, for some $x \in \mathbb{R}$.*

As with the case of no claim, when we pass to utilities defined on the entire real line the domain of the dual problem becomes simpler, being just the set $\mathcal{M}^a(S)$ (as opposed to the complicated set $\mathcal{D}$). In the same vein, the statement of the dual problem is much more transparent, since it does not involve the singular measures that we encountered before. It is simply (compare with (27))

$$v(y) = \inf_{Q \in \mathcal{M}^a(S)} E\left[V\left(y\frac{dQ}{dP}\right) - y\frac{dQ}{dP}B\right]. \tag{34}$$

The next theorem [22, theorem 1.1] provides the existence and uniqueness of solutions to the hedging problem for utilities defined on the entire $\mathbb{R}$. The remark following theorem 2.7 about the optimal measure $\widehat{Q}(y)$ being actually equivalent to P when $U(\infty) = \infty$ applies here as well (as can be seen from the form of the dual problem (34)).

Theorem 2.10 *Suppose that assumptions 1, 2 (case 2), 5 and 7 are satisfied. Then:*

1. *For any $x \in \mathbb{R}$ and $y > 0$, the problems*

$$u(x) = \sup_{f \in \mathcal{C}_U(x)} E[U(f - B)], \quad v(y) = \inf_{Q \in M^a(S)} E\left[V\left(y\frac{dQ}{dP}\right) - y\frac{dQ}{dP}B\right]$$

have unique optimizers $\widehat{f}(x) \in C_U(x)$ and $\widehat{Q}(y) \in \mathcal{M}^a(S)$ satisfying

$$U'(\widehat{f}(x) - B) = y\frac{d\widehat{Q}(y)}{dP}, \tag{35}$$

where x and y are related by $u'(x) = y$.

2. *If it occurs that $\widehat{Q}(y) \in \mathcal{M}^e(S)$, then $\widehat{f}(x)$ equals the terminal value $\widehat{X}_T(x)$ of a uniformly integrable $\widehat{Q}(y)$-martingale of the form*

$$\widehat{X}_t(x) = x + (\widehat{H}(x) \cdot S)_t,$$

for some $\widehat{H}(x) \in L(S)$.

To assert that the optimal measure $\widehat{Q}(y)$ is actually equivalent to P for the case of exponential utility, the analogue of proposition 2.8, we follow [10] and consider the change from P to an equivalent probability measure P_B with density

$$\frac{dP_B}{dP} = c_B e^{\gamma B}, \text{ with } c_B^{-1} = E[e^{\gamma B}]. \tag{36}$$

Therefore, for any $Q \ll P$, we have that

$$E\left[\frac{dQ}{dP}\log\frac{dQ}{dP}\right] = E_{P_B}\left[\frac{dQ}{dP_B}\log\frac{dQ}{dP_B}\right] + \log c_B + E\left[\gamma\frac{dQ}{dP}B\right]. \tag{37}$$

It then follows from the boundedness of B that Q has finite relative entropy with respect to P if and only if it has finite relative entropy with respect to P_B.

Now notice that the dual problem in this case is

$$\begin{aligned} v(y) &= \inf_{Q \in \mathcal{M}^a(S)} E\left[\frac{y}{\gamma}\frac{dQ}{dP}\left(\log\left(y\frac{dQ}{dP}\right) - 1\right) - y\frac{dQ}{dP}B\right] \\ &= \frac{y}{\gamma}(\log y - 1) + \inf_{Q \in \mathcal{M}^a(S)} E\left[\frac{dQ}{dP}\log\frac{dQ}{dP} - \gamma\frac{dQ}{dP}B\right] \\ &= \frac{y}{\gamma}(\log y c_B - 1) + \frac{y}{\gamma}\inf_{Q \in \mathcal{M}^a(S)} E_{P_B}\left[\frac{dQ}{dP_B}\log\frac{dQ}{dP_B}\right], \end{aligned} \tag{38}$$

from which we see that its minimizer coincides with the minimizer of the relative entropy with respect to P_B over all the absolutely continuous martingale measures for S. But from the argument preceding corollary 2.8, such a minimizer is equivalent to P_B (and therefore to P) provided there is at least one Q in $\mathcal{M}^e(S)$ whose relative entropy with respect to P_B is finite, which in turn is the same as having at least one Q in $\mathcal{M}^e(S)$ whose relative entropy with respect to P is finite. This suffices to prove:

Corollary 2.11 *Let $U(x) = -\frac{e^{-\gamma x}}{\gamma}$, $\gamma > 0$, and suppose that assumptions 1, 5 and 7 hold. If in addition we have that*

$$E\left[\frac{dQ}{dP}\log\left(\frac{dQ}{dP}\right)\right] < \infty, \tag{39}$$

for some $Q \in \mathcal{M}^e(S)$, then the minimizer $\widehat{Q}(y)$ of theorem 2.10 is the equivalent local martingale measure $\widehat{Q}$, independent of $y > 0$, which minimizes the relative entropy with respect to P_B among all absolutely continuous martingale measures. Therefore $\widehat{f}(x)$ equals the terminal value $\widehat{X}_T(x)$ of a uniformly integrable $\widehat{Q}$-martingale of the form

$$\widehat{X}_t(x) = x + (\widehat{H}(x) \cdot S)_t,$$

for some $\widehat{H}(x) \in L(S)$.

We end this review section with a discussion about complete markets, defined to be those for which there is exactly one equivalent martingale measure Q, that is, $\mathcal{M}^e(S)$ is the singleton $\{Q\}$. The second fundamental theorem of asset pricing relates this definition with the existence of a replicating portfolio for each bounded $\mathcal{F}_T$–measurable random variable.

Theorem 2.12 (FTAP II) *Suppose that assumption 1 holds. Then the following are equivalent:*

1. *The market is complete (i.e. $\mathcal{M}^e(S) = \{Q\}$).*
2. *For each $X \in L^\infty(\Omega, \mathcal{F}_T, P)$ there exist a unique admissible portfolio $H \in \mathcal{H}$ and a constant $x \in \mathbb{R}$ such that*

$$X = x + (H \cdot S)_T. \tag{40}$$

For complete markets, Merton's problem can be solved almost explicitly in terms of $\frac{dQ}{dP}$. The results of the next two theorems, which are slightly stronger versions of [20, theorem 2.0] and [26, theorem 2.1] (since we are assuming reasonable asymptotic elasticity for all our utility functions), are the analogues of theorems 2.6 and 2.7 for complete markets. Notice that for either case 1 or case 2 in assumption 2, the value function for the dual problem is reduced to

$$v(y) = E\left[V\left(y\frac{dQ}{dP}\right)\right], \quad y > 0 \tag{41}$$

(for case 1 this was proved in [20, lemma 4.3]; for case 2 it is trivial, since $\mathcal{M}^e(S) = \{Q\}$ implies that $\mathcal{M}^a(S) = \{Q\}$ as well).

Theorem 2.13 *Suppose that $\mathcal{M}^e(S) = \{Q\}$ and assumptions 2 (case 1) and 3 hold. Then, for any $x \in dom(U)$, the problem*

$$u(x) = \sup_{X_T \in C(x)} E[U(X_T)] \tag{42}$$

has a unique optimizer $\widehat{X}_T(x) \in C(x)$ given by

$$\widehat{X}_T(x) = -V'\left(y\frac{dQ}{dP}\right), \tag{43}$$

where y is the solution to the equation

$$E\left[-V'\left(y\frac{dQ}{dP}\right)\frac{dQ}{dP}\right] = x. \tag{44}$$

Theorem 2.14 *Suppose that $\mathcal{M}^e(S) = \{Q\}$ and assumptions 2 (case 2) and 4 hold. Then, for any $x \in \mathbb{R}$, the problem*

$$u(x) = \sup_{f \in C_U(x)} E[U(f)], \tag{45}$$

has a unique optimizer $\widehat{f}(x) \in C_U(x)$ given by

$$\widehat{f}(x) = -V'\left(y\frac{dQ}{dP}\right), \tag{46}$$

where y is the solution to the equation

$$E\left[-V'\left(y\frac{dQ}{dP}\right)\frac{dQ}{dP}\right] = x. \tag{47}$$

Moreover, $\widehat{f}(x)$ equals the terminal value $\widehat{X}_T(x)$ of a uniformly integrable Q–martingale of the form

$$\widehat{X}_t(x) = x + (\widehat{H}(x) \cdot S)_t, \quad t \in [0, T],$$

for some $\widehat{H}(x) \in L(S)$.

There is no need to state versions of theorems 2.9 and 2.10, since for complete markets the solution to the hedging problem (1) for a bounded claim B can be expressed in terms of the solution of Merton's problem. Indeed, by theorem 2.12, there exists (B_0, H^B) such that

$$B = B_0 + (H^B \cdot S)_T$$

and this can now be used to write (1) in the form of the Merton problem

$$\sup_{H \in \mathcal{H}} E\left[U\left(x - B_0 + ((H - H^B) \cdot S)_T\right)\right].$$

Therefore, if $\widehat{H}^0(x - B_0)$ is the optimal portfolio for the Merton problem starting with wealth $x - B_0$ obtained either from theorem 2.13 or from theorem 2.14, then the optimal portfolio for the hedging problem for the claim B starting with wealth x will be given by

$$\widehat{H}(x) = \widehat{H}^0(x + B_0) + H^B. \tag{48}$$

3 The dynamics of portfolio selection

The theorems of the previous section give precise statements ensuring the existence and uniqueness of solutions for both the optimal investment and optimal hedging problems for different types of utility functions. We have seen that under well defined conditions, there is a clear sense in which the optimal solution can always be approximated arbitrarily well by trading according to admissible portfolios. In what follows, to adopt a unified notation, we will write $H \in \mathcal{A}$, which loosely stands for "allowed" portfolios. In the back of our minds, however, we will keep the rigorous notion of what it stands for: admissible portfolios which, starting with initial capital $x \in \mathbb{R}$, generate terminal wealths in $C(x)$ and $\mathcal{C}(x)$ for theorems 2.6 and 2.9, respectively, or terminal wealths whose utilities arbitrarily approximate the utility of the optimal solutions (in the L^1 sense) $\widehat{f}(x)$ for theorems 2.7 and 2.10. We also use the notation $\mathcal{A}_{(s,t]}$ for portfolio processes defined only on the time interval $(s, t]$, as well as the shorthand for stochastic integration in this interval

$$(H \cdot S)_s^t := \int_s^t H_u dS_u, \quad 0 \leq s \leq t \leq T.$$

Consistently with our previous section we have that $\mathcal{A} = \mathcal{A}_{(0,T]}$ and $(H \cdot S)_t = (H \cdot S)_0^t$, $t \in [0, T]$.

To understand better the optimal selection problem it is useful to formulate a dynamical version of it. Let us write $\widehat{H}^{(x,0)}$ for the optimal solution to the static primal problem

$$u(x) = \sup_{H \in \mathcal{A}} E[U(x + (H \cdot S)_T - B)], \tag{49}$$

obtained according to the theorems of the previous section, that is, starting at time 0 with initial wealth x. For any intermediate time $t \in [0, T]$ and $x \in \operatorname{dom}(U)$, we can write

$$\begin{aligned} u(x) &= \sup_{H \in \mathcal{A}} E[U(x + (H \cdot S)_T - B)] \\ &= \sup_{H \in \mathcal{A}_{(0,t]}} E\left[\sup_{H \in \mathcal{A}_{(t,T]}} E_t[U(x + (H \cdot S)_0^t + (H \cdot S)_t^T - B)] \right], \end{aligned} \tag{50}$$

which leads us to the study of the *conditional* problem

$$u_t(\mathsf{w}) = \sup_{H \in \mathcal{A}_{(t,T]}} E_t[U(\mathsf{w} + (H \cdot S)_t^T - B)], \tag{51}$$

where $\mathsf{w} \in \mathbb{R}$ represents the wealth accumulated up to time t. If we trade according to $\widehat{H}^{(x,0)}$ up to time t, that is, if $\mathsf{w} = x + (\widehat{H}^{(x,0)} \cdot S)_t$, then we must have

$$u_t(\mathsf{w}) = E_t[U(\mathsf{w} + (\widehat{H}^{(\mathsf{w},t)} \cdot S)_t^T - B)], \tag{52}$$

for some portfolio $\widehat{H}^{(\mathsf{w},t)} \in \mathcal{A}_{(t,T]}$ (that is, starting at time t with wealth w) which agrees with the restriction of $\widehat{H}^{(x,0)}$ on the interval $(t, T]$. In other words, the optimal portfolio $\widehat{H}^{(x,0)}$ is also conditionally optimal. This is a special instance of the dynamic programming principle, which for this stochastic control problem has the form

$$u_s(\mathsf{w}) = \sup_{H \in \mathcal{A}_{(s,t]}} E_s[u_t(\mathsf{w} + (H \cdot S)_s^t)], \tag{53}$$

for $0 \leq s \leq t \leq T$.

The certainty equivalent value and the indifference price

There is a useful way to view the value function $u_t(\mathsf{w})$. By the intermediate value theorem, $U^{-1}(u_t(\mathsf{w}))$ exists for each (w, t), P–almost surely. This defines, for each (w, t), the random variable

$$B_t(\mathsf{w}) = \mathsf{w} - U^{-1}(u_t(\mathsf{w})), \tag{54}$$

which can be called the *certainty equivalent value* of the claim B at time t. Since

$$U(\mathsf{w} - B_t(\mathsf{w})) = E_t[U(\mathsf{w} + (\widehat{H}^{(\mathsf{w},t)} \cdot S)_t^T - B)],$$

the certain utility achieved by investing the amount $\mathsf{w} - B_t(\mathsf{w})$ in the risk free account equals the expected utility of the terminal wealth $\mathsf{w} + (\widehat{H}^{(\mathsf{w},t)} \cdot S)_t^T - B$ of the optimal hedging portfolio. For Merton's problem, where $B \equiv 0$, the amount $-B_t^0(\mathsf{w})$ indicates by how much the optimally invested portfolio outperforms the constant portfolio w over the period $(t, T]$. By putting $s = 0$ in (53) we obtain

$$\begin{aligned} u(x) &= \sup_{H \in \mathcal{A}_{(0,t]}} E[u_t(x + (H \cdot S)_0^t] \\ &= \sup_{H \in \mathcal{A}_{(0,t]}} E[U[x + (H \cdot S)_0^t - B_t(x + (H \cdot S)_0^t)]. \end{aligned} \tag{55}$$

Therefore, $B_t(\mathsf{w})$ represents a wealth dependent effective value of the claim B at time t.

Following [16] (according to [1]), a clear interpretation of the certainty equivalent values can be given by considering an investor who, holding wealth w at time t, 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 by holding the portfolio $\widehat{H}^{(\mathsf{w}+\pi,t)}$, he will achieve an expected utility

$$E_t[U(\mathsf{w} + \pi + (\widehat{H}^{(\mathsf{w}+\pi,t)} \cdot S)_t^T - B)] = U(\mathsf{w} + \pi - B_t(\mathsf{w} + \pi))$$

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

$$E_t[U(\mathsf{w} + (\widehat{H}^{(\mathsf{w},t)}(0) \cdot S)_t^T)] = U(\mathsf{w} - B_t^0(\mathsf{w})).$$

The *indifference price* of the claim B at time t for wealth w is the value for $\pi = \pi_t^B(\mathsf{w})$ which makes these equal, that is, it is the solution of

$$\pi_t^B(\mathsf{w}) = B_t(\mathsf{w} + \pi_t^B(\mathsf{w})) - B_t^0(\mathsf{w}). \tag{56}$$

Since we have defined these concepts from the point of view of an agent faced with a liability B at time t, this indifference price corresponds to a "seller's price". To obtain the correct notion of a "buyer's price", we just need to consider the reverse claim $-B$, which then produces a terminal wealth with expected utility equaling that of $\mathsf{w} - \pi - B_t(\mathsf{w} - \pi)$ when bought by π. The indifference price is now the value of π that makes this equal to the amount whose certain utility equals the expected utility for Merton's problem starting with wealth w at time t, which by definition is $\mathsf{w} - B_t^0(\mathsf{w})$. In other words, it is the solution of

$$\pi = B_t^0(\mathsf{w}) - B_t(\mathsf{w} - \pi), \tag{57}$$

which therefore equals $-\pi_t^B(\mathsf{w})$ as defined in (56).

In a complete market, the indifference price equals the risk-neutral price, that is, if the bounded claim B is written in terms of its unique replicating admissible portfolio H^B as $B = B_0 + (H^B \cdot S)_T$, then

$$\pi_t^B = B_0 + (H^B \cdot S)_0^t = E_{t,Q}[B], \tag{58}$$

where Q is the unique equivalent martingale measure. The first equality above remains true in incomplete markets if the claim B happens to satisfy $B = B_0 + (H^B \cdot S)_T$ for some admissible portfolio H^B.

The Davis price

Let us assume for a moment that the solutions of the dual problems in theorems 2.6 and 2.7 are equivalent martingale measures (in case 2 we have seen that this indeed the case for the exponential utility under the finite entropy condition; for counterexamples where in case 1 the solution fails to be a martingale, see [20]). If, for each $\varepsilon \geq 0$, we let $B_t^\varepsilon(\mathsf{w})$ denote the certainty equivalent value of εB, then the Davis price of B is defined to be [9]

$$\pi_t^{Davis}(\mathsf{w}) = \left.\frac{dB_t^\varepsilon(\mathsf{w})}{d\varepsilon}\right|_{\varepsilon=0}. \tag{59}$$

By differentiating the identity

$$U(\mathsf{w} - B_t^\varepsilon(\mathsf{w})) = E_t[U(\mathsf{w} + (\widehat{H}^{(\varepsilon,\mathsf{w},t)} \cdot S)_t^T - \varepsilon B)]$$

at $\varepsilon = 0$ and noting that, by optimality,

$$\left.\frac{d\widehat{H}^{(\varepsilon,\mathsf{w},t)}}{d\varepsilon}\right|_{\varepsilon=0} = 0,$$

we see that

$$\pi_t^{Davis}(\mathsf{w}) = \frac{E_t[U'(\mathsf{w} + (\widehat{H}^{(0,\mathsf{w},t)} \cdot S)_t^T)B]}{U'(\mathsf{w} - B_t^0(\mathsf{w}))}.$$

But from the theory of the Merton problem, a dynamical version of either (13) or (19) gives

$$U'(\mathsf{w} + (\widehat{H}^{(0,\mathsf{w},t)} \cdot S)_t^T) = U'(\mathsf{w} - B_t^0(\mathsf{w}))\frac{d\widehat{Q}_t(y)}{dP}$$

for $y = u_t'(\mathsf{w})$, where $\widehat{Q}_t(y)$ stands for the optimal solution to the conditional dual problem. Thus the Davis price of B is given by the expectation pricing

$$\pi_t^{Davis}(\mathsf{w}) = E_{t,\widehat{Q}_t(y)}[B]. \tag{60}$$

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.

Exponential utility

An important simplification occurs if we specialize to the exponential utility $U(x) = -\frac{e^{-\gamma x}}{\gamma}$, $\gamma > 0$. A look at (51) shows that $u_t(\mathsf{w})$ factorizes as

$$u_t(\mathsf{w}) = -\frac{e^{-\gamma \mathsf{w}}}{\gamma} \inf_{H \in \mathcal{A}_{(t,T]}} E_t \left[e^{-\gamma (H\cdot S)_t^T + \gamma B} \right] =: -\frac{e^{-\gamma \mathsf{w}}}{\gamma} v_t. \tag{61}$$

Here we see that v_t is a time dependent but wealth independent $\mathcal{F}_t$–measurable random variable. We also see that the certainty equivalent value

$$B_t = -\frac{1}{\gamma} \log v_t, \tag{62}$$

the optimal portfolio $\widehat{H}^{(t)}$ and the indifference price π_t are all wealth independent processes.

Example (geometric Brownian motion): Consider now a market of d stocks whose prices, discounted by the constant interest rate r, satisfy

$$\frac{dS_t^i}{S_t^i} = (\mu^i - r)dt + \sum_{\alpha=1}^{d} \sigma^{i\alpha} dW^\alpha, \tag{63}$$

where $\mu^i \in \mathbb{R}$ and the invertible $d \times d$ matrix $\sigma^{i\alpha}$ are constants and (W^α) is a d–dimensional P–Brownian motion. This market is complete and, as is well known, the unique equivalent martingale measure Q has Radon-Nikodym derivative

$$\frac{dQ}{dP} = \exp\left[-\int_0^T \left(\sum_\alpha \lambda^\alpha dW^\alpha + \frac{1}{2} \|\lambda\|^2 dt \right) \right], \tag{64}$$

with constant market price of risk $\lambda^\alpha = \sum_i (\sigma^{-1})^{\alpha i}(\mu^i - r)$.

For the exponential utility function with initial wealth x, the optimal discounted terminal wealth $\widehat{X}_T$ is given by

$$e^{-\gamma \widehat{X}_T} = y \frac{dQ}{dP}, \tag{65}$$

for a constant y to be determined. From this and (64), one finds

$$\widehat{X}_T = -\frac{1}{\gamma}\left(\log y + \frac{1}{2}\int_0^T \|\lambda\|^2 dt\right) + \frac{1}{\gamma}\int_0^T \sum_{ij}(\mu^i - r)\left((\sigma\sigma^T)^{-1}\right)^{ij}\frac{dS_t^j}{S_t^j},$$

so the optimal portfolio for Merton's problem is

$$(\widehat{H}^0)_t^j = \frac{\sum_i \left((\sigma\sigma^T)^{-1}\right)^{ij}(\mu^i - r)}{\gamma S_t^j} \tag{66}$$

and y is the solution to the equation

$$x = -\frac{1}{\gamma}\left(\log y + \frac{1}{2}\|\lambda\|^2 T\right). \tag{67}$$

The certainty equivalent value for Merton's problem in this market turns out to be

$$B_t^0 = \frac{1}{2}\|\lambda\|^2(t - T). \tag{68}$$

Since the market is complete, the indifference price of any bounded claim B equals its risk-neutral price (its Black-Scholes price), so the certainty equivalent value is given by

$$B_t = \frac{1}{2}\|\lambda\|^2(t - T) + E_{t,Q}[B]. \tag{69}$$

Finally, the optimal hedging portfolio for the claim B is

$$\widehat{H}_t = \widehat{H}_t^0 + H_t^B, \tag{70}$$

where H_t^B is the replicating portfolio for B.

4 Discrete time hedging

We now restrict to discrete time hedging, where the portfolio processes have the form

$$H_t = \sum_{k=1}^{K} H_k \, \mathbf{1}_{(t_{k-1},t_k]}(t) \tag{71}$$

where each H_k is an $\mathbb{R}^d$–valued $\mathcal{F}_{k-1}$ random variable. We take the discrete time partition of the interval $[0, T]$ to be of the form

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

and use the notation $S_j := S_{t_j}$ for discrete time stochastic processes. The discounted wealth process will be $X_j = x + (H \cdot S)_j$, with the notation $(H \cdot S)_j := \sum_{k=1}^{j} H_k \Delta S_k$, $(H \cdot S)_k^j := (H \cdot S)_j - (H \cdot S)_k$ and $\Delta S_k := S_k - S_{k-1}$.

Now the dynamic programming problem (53) falls into K subproblems

$$u_{k-1}(x) = \sup_{H \in \mathcal{A}_{(t_{k-1}, t_k]}} E_{k-1}[u_k(x + H_k \Delta S_k)], \qquad k = K, K-1, \ldots, 1 \quad (72)$$

subject to the terminal condition $u_K(x) = U(x)$. Then, for each x this defines a process $u_k(x)$. Similarly, the certainty equivalent value process $B_k(x)$ is defined iteratively by

$$U(x - B_{k-1}(x)) = \sup_{H \in \mathcal{A}_{(t_{k-1}, t_k]}} E_{k-1}[U(x + H_k \Delta S_k - B_k(x + H_k \Delta S_k)] \quad (73)$$

with $B_K(x)$ taken equal to the terminal claim B. In both formulations of the problem, let $\widehat{H}_k(x)$ denote the minimizer, which of course is an $\mathcal{F}_{k-1}$ random variable.

In what follows, we will be largely concerned with markets and claims which satisfy the Markovian conditions:

Assumption 8 *The market is Markovian and its state variables* $Z = (S^1, \ldots, S^d, Y^1, \ldots, Y^{n-d})$ *lie in a finite dimensional state space* $\mathcal{S} \in \mathbb{R}^n$.

Assumption 9 *The contingent claim is taken to be of the form* $B_T = \Phi(Z_T)$ *for a bounded Borel function* $\Phi : \mathcal{S} \to \mathbb{R}$.

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. In this Markovian setting, the solution of (72) and the optimal allocation have the form

$$u_k(x) = g_k(x, Z_k) \quad (74)$$

$$\widehat{H}_{k+1}(x) = h_{k+1}(x, Z_k) \quad (75)$$

for (deterministic) Borel functions $\{g_k, h_{k+1}\}_{k=0}^{K-1}$ mapping $\mathrm{dom}(U) \times \mathcal{S}$ to $\mathbb{R}$ and $\mathbb{R}^d$ respectively. Similarly, the solution of (73) has the optimal allocation $\widehat{H}_{k+1}$ as above and B_k in the form

$$B_k(x) = b_k(x, Z_k) \quad (76)$$

for Borel functions $\{b_k\}_{k=0}^{K-1}$ mapping $\mathrm{dom}(U) \times \mathcal{S}$ to $\mathbb{R}$.

As indicated in the previous section, matters simplify in the special case of the exponential utility function $U(x) = -\frac{e^{-\gamma x}}{\gamma}$, $\gamma > 0$. One finds that the dynamic program can be written in the wealth independent form $u_k(x) = U(x)v_k$, $\widehat{H}_k(x) = \widehat{H}_k$, and $B_k(x) = B_k$ where the random variables v, $\widehat{H}$ and B have the form:

$$\begin{aligned} v_k &= g_k(Z_k) && (77)\\ \widehat{H}_{k+1} &= h_{k+1}(Z_k) && (78)\\ B_k &= b_k(Z_k) && (79) \end{aligned}$$

for deterministic functions g_k, h_{k+1}, and b_k on the state space $\mathcal{S}$. The iteration equations are simply

$$\begin{aligned} g_k(Z) &= \inf_{h\in\mathbb{R}^d} E_k[\exp(-\gamma h \cdot \Delta S_{k+1}) g_{k+1}(Z_{k+1})|Z_k = Z] && (80)\\ \exp(\gamma b_k(Z)) &= \inf_{h\in\mathbb{R}^d} E_k[\exp(-\gamma(h \cdot \Delta S_{k+1} - b_{k+1}(Z_{k+1}))|Z_k = Z] && (81) \end{aligned}$$

and the optimal h defines the function $h_{k+1}(Z)$.

5 The exponential utility allocation algorithm

In this section we introduce a Monte Carlo method for learning the optimal trading strategy (75) for the discrete time Markovian problems discussed in the previous section. We want an algorithm which will generate an approximate trading rule, based on a data set $\{Z_k^i\}_{i=1,\dots,N;k=0,\dots,K}$ where $Z_k^i \in \mathbb{R}^n$ denotes the "state " of the ith sample path at time $t_k = kT/K$.

Consider the discrete time problem (72) for a general utility function U satisfying assumption 2. The optimal portfolio $\widehat{H}_{k+1}^i \in \mathbb{R}^d$ should be selected as $h_{k+1}(X_k^i, Z_k^i)$ 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 h_{k+1} from the data $\{Z\}$ will require being able to fill in the optimal wealth from time t_0 to t_k. Then, conditionally upon knowing the wealth X_k^i at time t_k, finding the function h_{k+1} requires dynamic programming backwards from time $t_K = T$ to t_k. In other words, a Monte Carlo learning algorithm for $\widehat{H}_k^i$ based on general utility will require both forward and backward dynamic programming. We have no effective method to suggest for general utility.

By contrast, in the special case of exponential utility, the theoretical optimal rule $\widehat{H}_{k+1}^i = h_j(Z_k^i)$ depends only on the directly observed data $\{Z_k^i\}$ and is independent of the wealth X_k^i. For this reason our algorithm works only for exponential utility, and we take

$$U(x) = -e^{-x},$$

for simplicity.

5.1 The algorithm

1. Step $k = K$: The final optimal allocation $\widehat{H}_K$ is defined to be the $\mathbb{R}^d$–valued $\mathcal{F}_{K-1}$ random variable which solves

$$\min_{H \in \mathcal{A}_{(K-1,K]}} E_{K-1}[\exp(-H \cdot \Delta S_K + B)], \qquad \text{(a.s.)}$$

which is easily seen to be equivalent to the optimizer of

$$\min_{H \in \mathcal{A}_{(K-1,K]}} E[\exp(-H \cdot \Delta S_K + B)] \tag{82}$$

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

$$\min_{h \in \mathcal{C}(\mathcal{S})} \Psi_K(h) \tag{83}$$

where $\Psi_K(h) := E[\exp(-h(Z_{K-1}) \cdot \Delta S_K + B)]$. On a finite set of data, we can pick an R–dimensional subspace $\mathcal{R}(\mathcal{S}) \subset \mathcal{B}(\mathcal{S})$ of functions on $\mathcal{S}$ and attempt to "learn" a suboptimal solution

$$h_K^{\mathcal{R}} = \arg\min_{h \in \mathcal{R}(\mathcal{S})} \Psi_K(h)$$

By the central limit theorem, the expectation $\Psi_{K-1}(h)$ for h in a neighbourhood of $h_K^{\mathcal{R}}$, and hence the solution $h_K^{\mathcal{R}}$ itself, can be approximated by the finite sample estimate

$$\widetilde{\Psi}_K(h) = \frac{1}{N} \sum_{i=1}^{N} \exp\left(-h(Z_{K-1}^i) \cdot \Delta S_K^i + \Phi(Z_K^i)\right) \tag{84}$$

This leads to the estimator $\widetilde{h}_K^{\mathcal{R}}$ based on $\{Z_k^i\}$ and the choice of subspace $\mathcal{R}$ defined by

$$\widetilde{h}_K^{\mathcal{R}} = \arg\min_{h \in \mathcal{R}(\mathcal{S})} \widetilde{\Psi}_K(h) \tag{85}$$

2. Inductive step for $k = K-1, \ldots, 2$: The estimate $\widetilde{h}_k^{\mathcal{R}}$ of the optimal rule $\widehat{h}_k$, for $2 \le k < K-1$ is determined inductively given the estimates $\widetilde{h}_{k+1}^{\mathcal{R}}, \ldots, \widetilde{h}_K^{\mathcal{R}}$. It is defined to be

$$\widetilde{h}_k^{\mathcal{R}} = \arg\min_{h \in \mathcal{R}(\mathcal{S})} \widetilde{\Psi}_k(h; \widetilde{h}_{k+1}^{\mathcal{R}}, \ldots, \widetilde{h}_K^{\mathcal{R}}) \tag{86}$$

where

$$\widetilde{\Psi}_k(h;\widetilde{h}^{\mathcal{R}}_{k+1},\ldots,\widetilde{h}^{\mathcal{R}}_K) =$$
$$\frac{1}{N}\sum_{i=1}^{N}\exp\left(-h(Z^i_k)\cdot\Delta S^i_{k+1}-\sum_{j=k+1}^{K}\widetilde{h}^{\mathcal{R}}_j(Z^i_j)\cdot\Delta S^i_j+\Phi(Z^i_K)\right) \quad (87)$$

3. Final step $k=1$: This step is degenerate since the initial values Z_0 are constant over the sample. Therefore we determine the optimal constant vector $\widetilde{h}_1\in\mathbb{R}^d$ by solving

$$\widetilde{h}_1=\arg\min_{h\in\mathbb{R}^d}\widetilde{\Psi}_1(h;\widetilde{h}^{\mathcal{R}}_2,\ldots,\widetilde{h}^{\mathcal{R}}_K) \quad (88)$$

To summarize, the algorithm above learns a collection of functions of the form $(\widetilde{h}_1,\widetilde{h}^{\mathcal{R}}_2,\ldots,\widetilde{h}^{\mathcal{R}}_K)\in\mathbb{R}^d\times\mathcal{R}(\mathcal{S})^{K-1}$ from the Monte Carlo simulation. This collection defines a suboptimal allocation strategy for the exponential hedging problem. Finally, the optimal value $\widetilde{\Psi}_1(\widetilde{h}_1;\widetilde{h}^{\mathcal{R}}_2,\ldots,\widetilde{h}^{\mathcal{R}}_K)$ is an estimate of the quantity e^{B_0}, where B_0 is the certainty equivalent value of the claim B at time $t=0$.

5.2 Discussion of errors

It is important to identify two distinct systematic sources of error in the algorithm. The first, which we call *approximation one*, is in focusing on suboptimal solutions $h^{\mathcal{R}}_k$ 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 $f_1,\ldots,f_R$ 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 indicate 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 [21] 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; the following is a heuristic argument to give a flavour of the problem for the kth step, $k \leq K$. By the central limit theorem, for a given confidence level $1-\alpha, \alpha \ll 1$, there exist constants C_1, C_2 so that

$$\|\Psi(h) - \widetilde{\Psi}(h)\| \leq \frac{C_1}{\sqrt{N}}, \quad \|\nabla\Psi(h) - \nabla\widetilde{\Psi}(h)\| \leq \frac{C_2}{\sqrt{N}}$$

with probability $1-\alpha$, for h in a convex neighbourhood of the true critical point $h^{\mathcal{R}}$, defined by $\nabla\Psi(h^{\mathcal{R}}) = 0$. We suppose that the estimated critical point $\widetilde{h}^{\mathcal{R}}$, defined by $\nabla\widetilde{\Psi}(\widetilde{h}^{\mathcal{R}}) = 0$, lies in this neighbourhood, and furthermore the operator inequalities $0 < C_3 \leq \nabla^2\Psi \leq C_4$ hold on the same neighbourhood. Then one immediately derives the inequalities

$$\|h^{\mathcal{R}} - \widetilde{h}^{\mathcal{R}}\| \leq \frac{C_2}{C_3\sqrt{N}} \tag{89}$$

$$|\Psi(h) - \widetilde{\Psi}(\widetilde{h}^{\mathcal{R}})| \leq \frac{C_1}{\sqrt{N}} + \frac{C_2^2 C_4}{2C_3^2 N} \tag{90}$$

which show convergence of $\widetilde{h}^{\mathcal{R}}$ to $h^{\mathcal{R}}$ for $N \to \infty$.

The above discussion addresses the errors made at the kth 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.

6 Numerical implementation

We have tested the algorithm in the simple problem of investment with exponential utility $U(x) = -e^{-x}$ in a stock which behaves as a geometric Brownian motion, with and without the purchase of a single at–the–money European put option. We have seen there is an exact solution to this problem which can be compared in detail to the solution generated by the Monte Carlo algorithm.

We consider the model of (63) with $d = 1$ and parameters $S_0 = 1$, $\mu = 0.1$, $\sigma = 0.2$ and $r = 0.0$ over the period of one year $T = 1$. We apply the allocation algorithm 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 buyer of a single written at–the–money European put. In each case we apply the method for simulations of length

$N = 1000, 10000, 100000$. Then, for comparison to theory, we use the same Monte Carlo simulations, but rehedged weekly according to the theoretical formula (70), with $-H^B$ equal to the Black-Scholes delta of the option.

Our results are displayed in figures 1 to 5. Figures 1,2,3,4 show the profit/loss distributions at time T for the learned Merton, learned put option, true Merton and true put option cases respectively. They show the empirical distributions for $N = 1000, 10000, 100000$. Figure 5 shows the values of the hedge ratio along a single sample path calculated according to both the strategy learned with $N = 100000$ and the true strategy.

For comparison of their performances, we tabulate below 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 $\gamma = 1/4$ (U_1), $\gamma = 1$ (U_2) and $\gamma = 4$ (U_3), 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.

From the U_2 values on the table, one can derive the learned estimates of the indifference price $0.0767, 0.0790$ and 0.0792 for the cases $N = 1000, 10000$ and 100000. Using the true strategy leads to the values $0.0798, 0.0796, 0.0795$, respectively. The theoretical Black-Scholes price is 0.0797.

Case	Mean	St. Dev.	U_1	U_2	U_3	VaR_{99}	VaR_{90}	CVaR_{99}	CVaR_{90}
1a	0.3572	0.5778	-0.9241	-0.8262	-3.3023	-1.0829	-0.3659	-1.2053	-0.6495
1b	0.2768	0.5136	-0.9409	-0.8674	-3.5224	-1.0159	-0.3743	-1.2144	-0.6543
1c	0.2528	0.5013	-0.9462	-0.8810	-2.7684	-0.9209	-0.3913	-1.0913	-0.6318
2a	0.4349	0.5797	-0.9064	-0.7652	-2.4748	-0.9828	-0.3012	-1.1348	-0.5720
2b	0.3562	0.5142	-0.9224	-0.8015	-2.9355	-0.9518	-0.2732	-1.1756	-0.5626
2c	0.3325	0.5020	-0.9275	-0.8139	-2.4439	-0.8532	-0.2859	-1.0633	-0.5387
3a	0.2307	0.4898	-0.9511	-0.8956	-2.5283	-0.9723	-0.3961	-1.0318	-0.6430
3b	0.2524	0.4945	-0.9460	-0.8773	-2.4184	-0.8852	-0.3778	-1.0429	-0.6096
3c	0.2506	0.4995	-0.9466	-0.8816	-2.6461	-0.9081	-0.3896	-1.0733	-0.6254
4a	0.3108	0.4904	-0.9322	-0.8269	-1.8302	-0.8878	-0.3135	-0.9492	-0.5628
4b	0.3322	0.4954	-0.9274	-0.8102	-1.7521	-0.8054	-0.2979	-0.9616	-0.5289
4c	0.3304	0.5005	-0.9279	-0.8142	-1.9178	-0.8274	-0.3098	-0.9921	-0.5449

Table 1: Mean, standard deviation, final expected utilities and risk measures for the profit/loss distribution of the learned Merton, learned put option, true Merton and true put option portfolios with (a) 1000, (b) 10000 and (c) 100000 Monte Carlo simulations of stock prices following a geometric Brownian motion.

7 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 [21]. 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 5, 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.

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 arises by restricting possible hedge strategies to a low dimensional subspace. It is clear that such a restriction will often lead to unsuitable strategies. However, we feel that approximation two, the finite sample size error, will likely be even more problematic for practical realizations. A brief study of the size of the constants which enter the estimates (89) and (90) suggests that reliable learned strategies will demand a very large value of N. In our simulations, $N = 100000$ gave reliable prices, but not hedging strategies. A third 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.

To conclude this discussion, it is worthwhile to revisit the way in which our method of dynamic programming (finding $\widehat{H}$ 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 $\widehat{H}$ 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.

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] D. Becherer. *Rational Hedging and Valuation with Utility-based Preferences.* PhD thesis, Technischen Universitat Berlin, 2001.

[2] F. Bellini and M. Frittelli. On the existence of minimax martingale measures. *Math. Finance*, 12(1):1–21, 2002.

[3] 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.

[4] W. Brannath and W. Schachermayer. A bipolar theorem for $L^0_+(\Omega, \mathcal{F}, P)$. In *Séminaire de Probabilités, XXXIII*, volume 1709 of *Lecture Notes in Math.*, pages 349–354. Springer, Berlin, 1999.

[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] I. Csiszár. I-divergence geometry of probability distributions and minimization problems. *Ann. Probability*, 3:146–158, 1975.

[7] J. Cvitanic, L. Goukasian, and F. Zapatero. Monte Carlo computation of optimal portfolios in complete markets. Preprint, 2000.

[8] J. Cvitanić, W. Schachermayer, and H. Wang. Utility maximization in incomplete markets with random endowment. *Finance Stoch.*, 5(2):259–272, 2001.

[9] 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.

[10] 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.

[11] F. Delbaen and W. Schachermayer. A general version of the fundamental theorem of asset pricing. *Math. Ann.*, 300(3):463–520, 1994.

[12] F. Delbaen and W. Schachermayer. The fundamental theorem of asset pricing for unbounded stochastic processes. *Math. Ann.*, 312(2):215–250, 1998.

[13] J. Detemple, R. Garcia, and M. Rindisbacher. A Monte Carlo method for optimal portfolios. To appear in Journal of Finance, 2003.

[14] M. Frittelli. Optimal solutions to utility maximization and to the dual problem. Preprint, 2002.

[15] 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.

[16] S. D. Hodges and A. Neuberger. Optimal replication of contingent claims under transaction costs. *Rev. Fut. Markets*, 8:222–239, 1989.

[17] Y. M. Kabanov and C. Stricker. On the optimal portfolio for the exponential utility maximization: remarks to the six-author paper. *Math. Finance*, 12(2):125–134, 2002.

[18] 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.

[19] I. Karatzas and S. E. Shreve. *Methods of mathematical finance*, volume 39 of *Applications of Mathematics*. Springer-Verlag, New York, 1998.

[20] 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.

[21] F. A. Longstaff and E. S. Schwartz. Valuing american options by simulation: a simple least-square approach. *Rev. Fin. Studies*, 14:113–147, 2001.

[22] M. P. Owen. Utility based optimal hedging in incomplete markets. *Ann. Appl. Probab.*, 12(2):691–709, 2002.

[23] P. Protter. *Stochastic integration and differential equations*, volume 21 of *Applications of Mathematics*. Springer-Verlag, Berlin, 1990. A new approach.

[24] R. T. Rockafellar. *Convex analysis*. Princeton Landmarks in Mathematics. Princeton University Press, Princeton, NJ, 1997. Reprint of the 1970 original, Princeton Paperbacks.

[25] W. Schachermayer. Optimal investment in incomplete financial markets. In H. Geman et al., editors, *Mathematical Finance: Bachelier Congress 2000*, pages 427–462, 2001.

[26] W. Schachermayer. Optimal investment in incomplete markets when wealth may become negative. *Ann. Appl. Probab.*, 11(3):694–734, 2001.

[27] W. Schachermayer. A supermartingale property of the optimal portfolio process. Submitted, 2002.

[28] V. Young and T. Zariphopoulou. Pricing dynamic insurance risks using the principle of equivalent utility. *Scand. Actuarial J.*, 4:246–279, 2002.

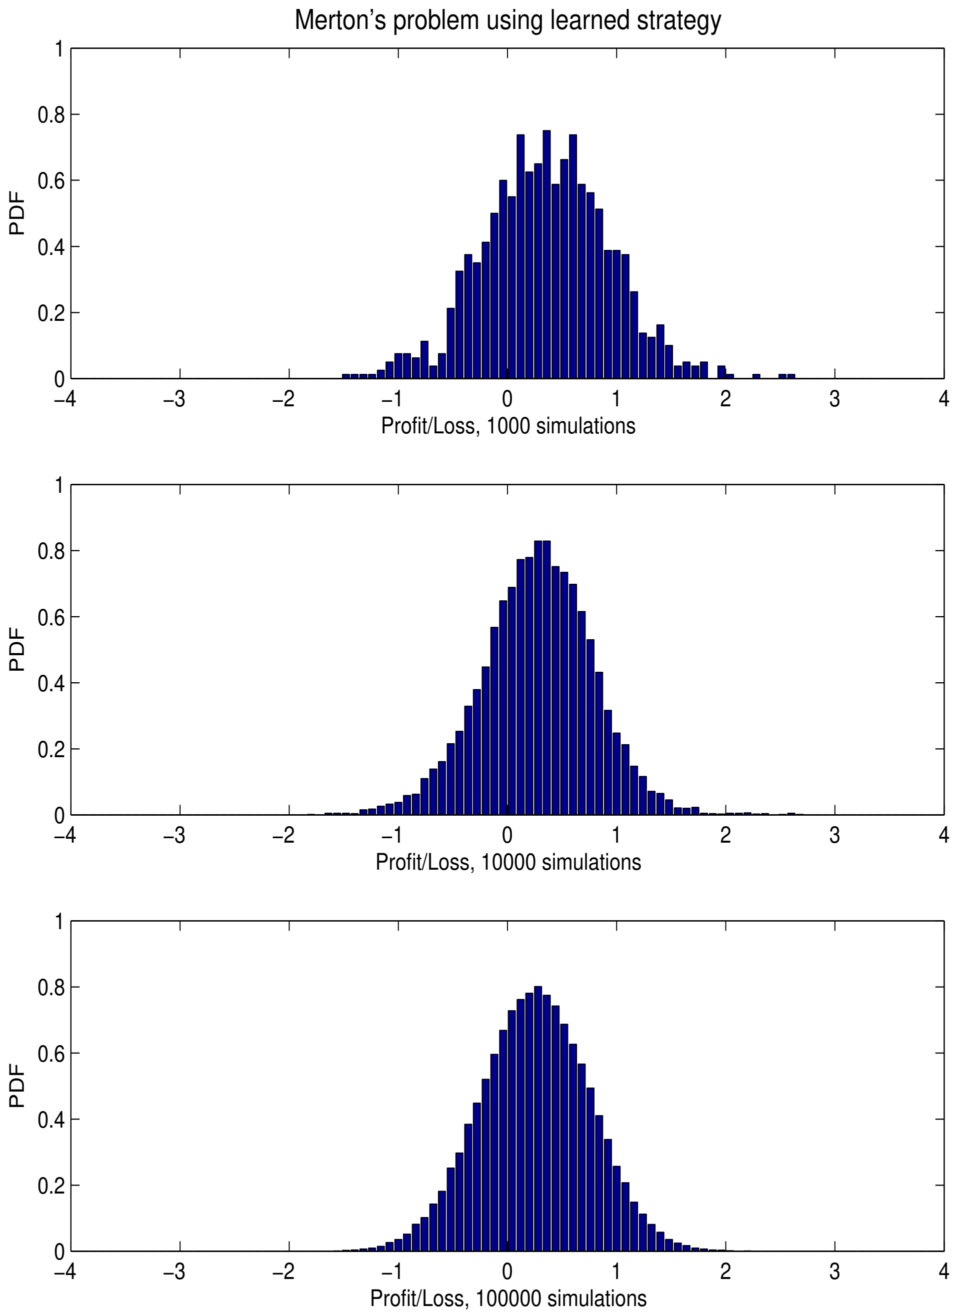

Figure 1: The profit/loss distribution of the learned investment portfolio, obtained from the exponential utility allocation algorithm as an approximated solution to Merton's problem, evaluated on simulated stock prices following a geometric Brownian motion.

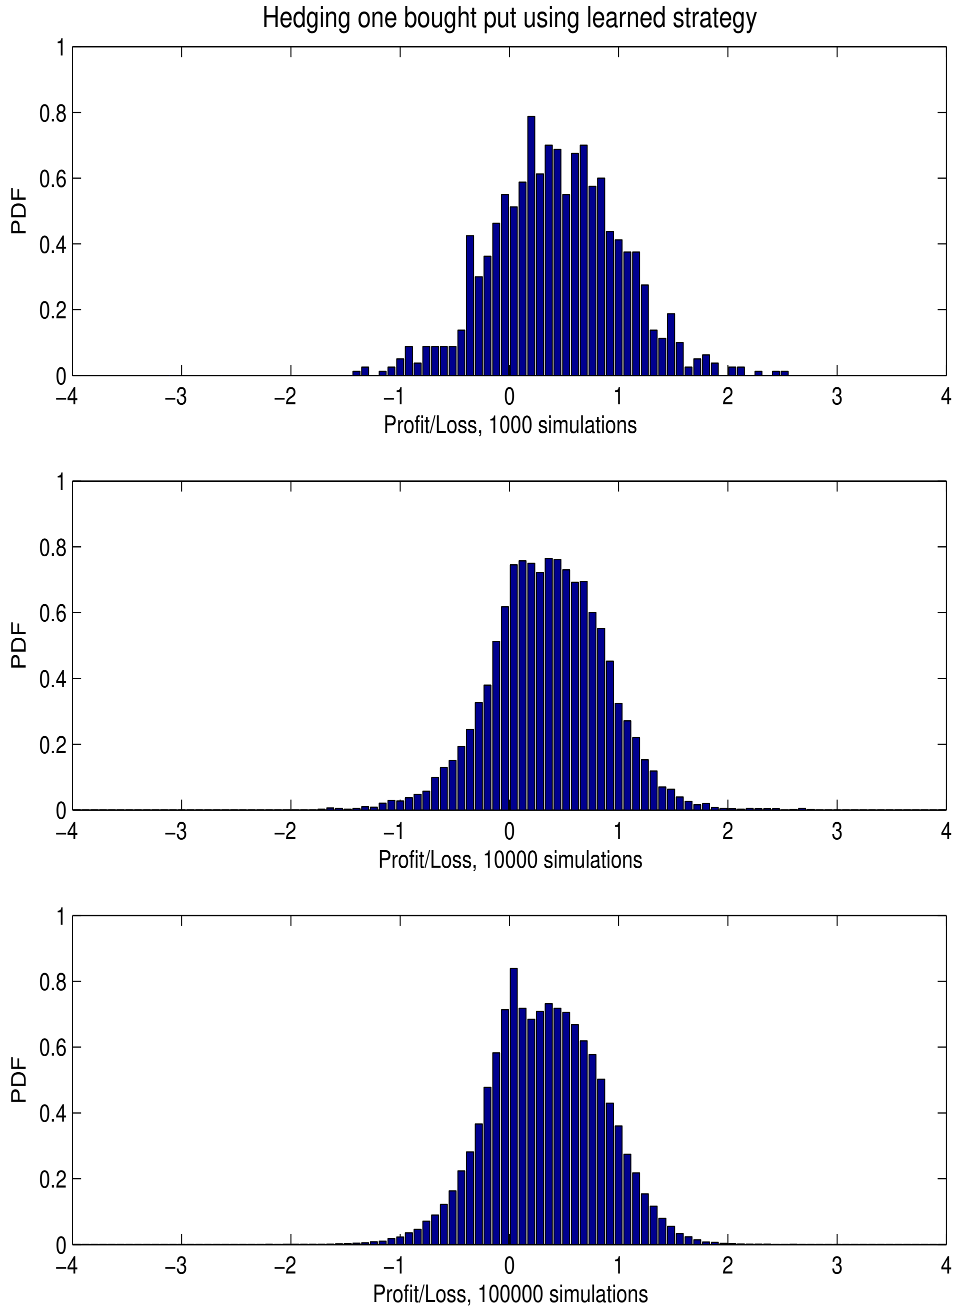

Figure 2: The profit/loss distribution of the learned hedging portfolio for the buyer of one put option, obtained from the exponential utility allocation algorithm on simulated stock prices following a geometric Brownian motion.

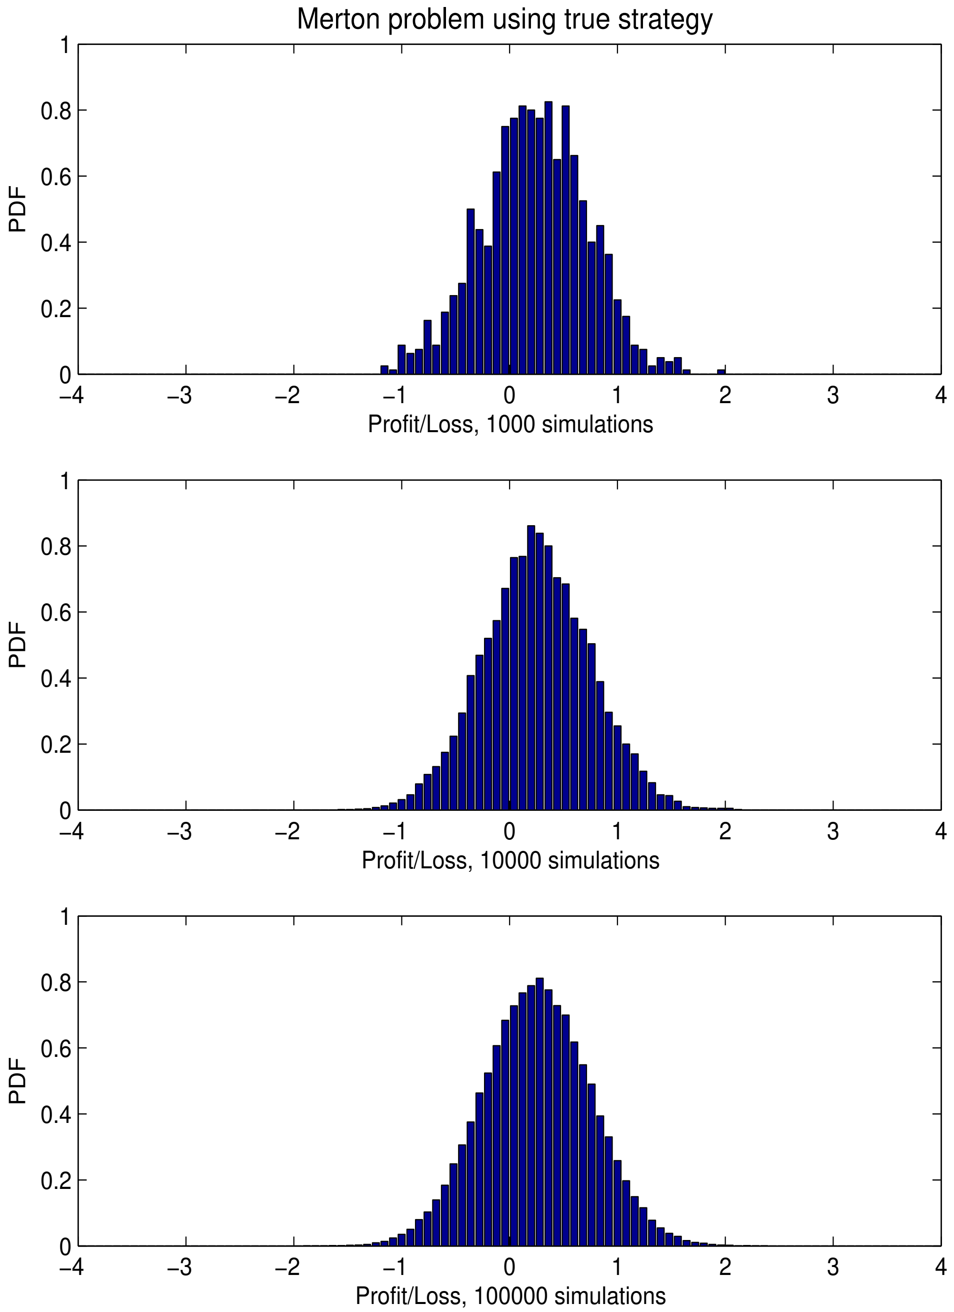

Figure 3: The profit/loss distributions of optimal investment portfolio, obtained as the exact solution for Merton's problem with exponential utility, evaluated on simulated stock prices following a geometric Brownian motion.

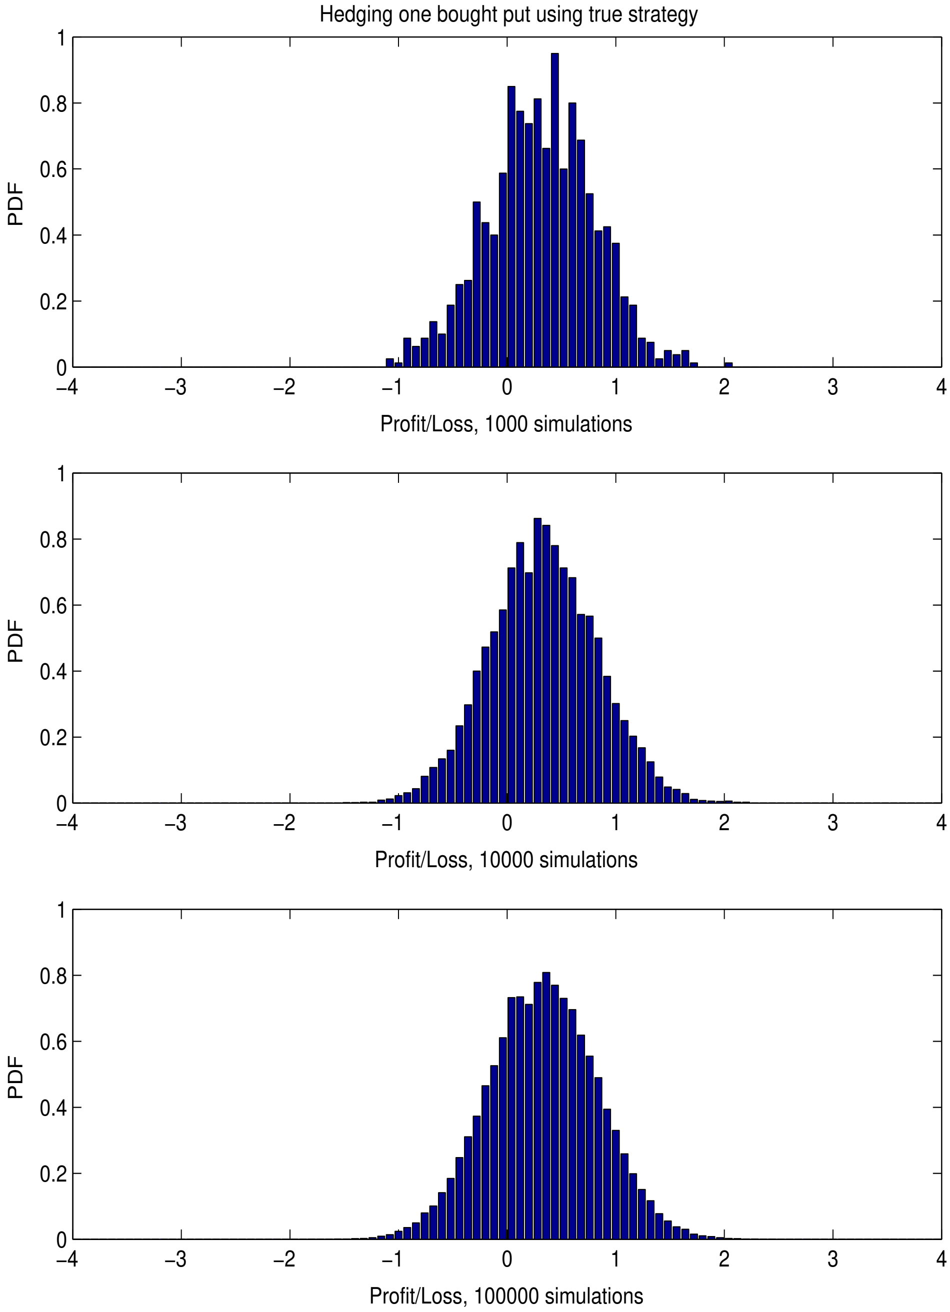

Figure 4: The profit/loss distributions of the optimal hedging portfolio for the buyer of one put option, obtained from Black-Scholes delta hedging combined with Merton's problem with exponential utility, evaluated on simulated stock prices following a geometric Brownian motion.

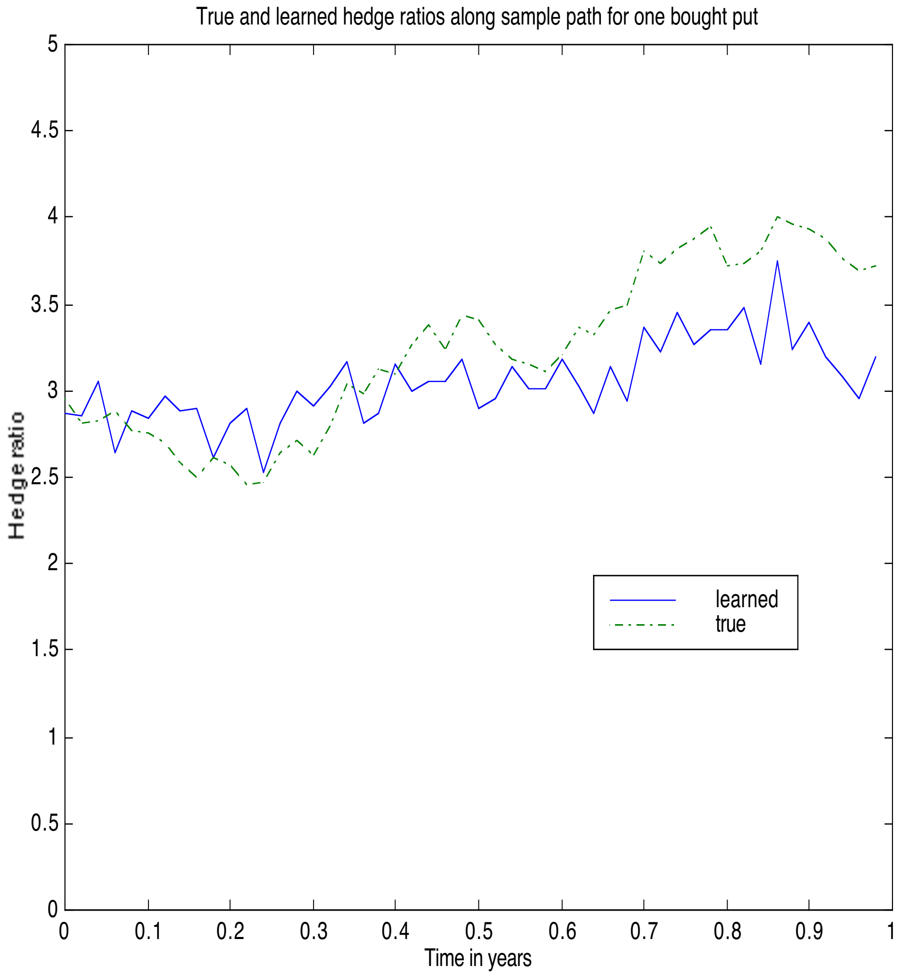

Figure 5: The hedge ratio (number of shares held) for the buyer of one put option on a simulated sample path of duration one year, for which the option matures in–the–money. The solid line shows the strategy learned with N=100000; the broken line shows the theoretical Black-Scholes-Merton strategy.