

Duality Theory and Simulation in Financial Engineering

Martin Haugh

Department of IE and OR, Columbia University, New York, NY 10027, martin.haugh@columbia.edu.

Abstract

This paper presents a brief introduction to the use of duality theory and simulation in financial engineering. It focuses on American option pricing and portfolio optimization problems when the underlying state space is high-dimensional. In general, it is not possible to solve these problems exactly due to the so-called “curse of dimensionality” and as a result, approximate solution techniques are required. Approximate dynamic programming (ADP) and dual based methods have recently been proposed for constructing and evaluating good approximate solutions to these problems. In this paper we describe these ADP and dual-based methods, and the role simulation plays in each of them. Some directions for future research are also outlined.

Keywords: Duality Theory, Simulation, Financial engineering.

1 INTRODUCTION

Portfolio optimization and American option pricing problems are among the most important problems in financial engineering. Portfolio optimization problems occur throughout the financial services as pension funds, mutual funds, insurance companies, endowments and other financial entities all face the fundamental problem of dynamically allocating their resources across different securities in order to achieve a particular goal. These problems are often very complex owing to their dynamic and stochastic nature, their high dimensionality and the complexity of real-world constraints. While researchers have developed very sophisticated models for addressing these problems, the current state-of-the-art is such that explicit solutions are available only in very special circumstances. (See, for example, Merton 1990, Cox and Huang 1991, Karatzas and Shreve 1997, and Liu 1998).

American option pricing has also presented several challenges to the financial engineering community. Even in the simple Black-Scholes framework (Black and Scholes 1973), a closed form expression for the price of an American put option is not available and so it must therefore be computed numerically. As pricing an American option amounts to solving an optimal stopping problem, Bellman's curse of dimensionality implies that pricing high-dimensional American options using standard numerical techniques is not practically feasible. Unfortunately, the same conclusion also applies to solving general high-dimensional portfolio optimization problems.

Because these high-dimensional problems occur frequently in practice, they are of considerable interest to both researchers and practitioners. In recent years there has been some success in tackling these problems using approximate dynamic programming (ADP) and dual-based methods. ADP methods (see, for example, Bertsekas and Tsitsiklis 1996) have had considerable success in tackling large-scale complex problems and have recently been applied successfully to problems in financial engineering (Brandt et al. 2001; Longstaff and Schwartz 2001; and Tsitsiklis and Van Roy 2001). One difficulty with ADP, however, is in establishing how far the sub-optimal ADP solution to a given problem is from optimality. In the context of optimal stopping problems and pricing American options, Haugh and Kogan (2001) and Rogers (2002) showed how a stochastic duality theory could be used to evaluate sub-optimal strategies, including those obtained from ADP methods. A stochastic duality theory also exists for portfolio optimization problems and this has been developed by many researchers in recent years (see, for example, Shreve and Xu 1992a and 1992b, He and Pearson 1991, Cvitanic and Karatzas 1992, and Karatzas and Shreve 1997). While this theory has had considerable success in characterizing optimal solutions, explicit solutions are still rare (see Rogers 2003). Recently Haugh, Kogan and Wang (2003) have shown how some of these dual formulations can be used to evaluate suboptimal policies by constructing lower and upper bounds on the true optimal value function. These suboptimal policies could be simple heuristic policies or policies resulting from some approximation techniques such as ADP.

Simulation techniques play a key role in both the ADP and dual-based evaluation methods that have been used to construct and evaluate solutions to these problems. While it has long been recognized that simulation is an indispensable tool for financial engineering (see the surveys of Boyle, Broadie and Glasserman 1997, and Staum 2002), it is only recently that simulation has begun to play an important role in solving *control* problems in financial engineering. These control problems include portfolio optimization and the pricing of American options, and they are the focus of this paper.

The remainder of the paper is outlined as follows. Sections 2 and 3 describe the American option pricing and portfolio optimization problems, respectively. We very briefly describe the ADP methods in Section 2.1 after which we will focus on the duality theory for both optimal stopping and portfolio optimization. Section 4 concludes and outlines some future research directions. Results will not be

presented in their full generality, and technical details will be often be omitted as we choose to focus instead on the underlying concepts and intuition.

2 PRICING AMERICAN OPTIONS

The Financial Market: We assume there exists a dynamically complete financial market that is driven by a vector-valued Markov process, $X_t = (X_t^1, \dots, X_t^n)$. In words, we say a financial market is dynamically complete if any random variable, W_T , representing a terminal cash-flow can be attained by using a *self-financing trading strategy*. (A self-financing trading strategy is a strategy where changes in the value of the portfolio are only due to capital gains or losses. In particular, no net addition or withdrawal of funds is allowed after date $t = 0$ and any new purchases of securities must be financed by the sale of other securities.) X_t represents the time t vector of risky asset prices as well as the values of any relevant state variables in the market. We also assume there exists a risk-free security whose time t price is $B_t = e^{rt}$, where r is the continuously compounded risk-free rate of interest. Finally, since markets are assumed to be dynamically complete, there exists (see Duffie 1996) a unique risk-neutral valuation measure, \mathcal{Q} .

Option Payoff: Let $h_t = h(X_t)$ be a nonnegative adapted process representing the payoff of the option so that if it is exercised at time t the holder of the option will then receive h_t .

Exercise Dates: The American feature of the option allows the holder of the option to *exercise* it at any of the pre-specified exercise dates in $\mathcal{T} = \{0, 1, \dots, T\}$.

Option Price: The value process of the American option, V_t , is the price process of the option conditional on it not having been exercised before t . It satisfies

$$V_t = \sup_{\tau \geq t} \mathbb{E}_t^{\mathcal{Q}} \left[\frac{B_t h_\tau}{B_\tau} \right]. \quad (1)$$

where τ is any stopping time with values in the set $\mathcal{T} \cap [t, T]$.

If X_t is high-dimensional, then standard solution techniques such as dynamic programming become impractical and we cannot hope to solve the optimal stopping problem (1) exactly. Fortunately, efficient ADP algorithms for addressing this problem have recently been developed independently by Longstaff and Schwartz (2001) and Tsitsiklis and Van Roy (2001). We now briefly describe the main ideas behind these algorithms, both of which rely on the ability to simulate paths of the underlying state vectors.

2.1 ADP for Pricing American Options

Once again, the pricing problem at time $t = 0$ is to compute

$$V_0 = \sup_{\tau \in \mathcal{T}} E_0^{\mathcal{Q}} \left[\frac{h_\tau}{B_\tau} \right]$$

and in theory this problem is easily solved using value iteration. In particular, we would obtain

$$\begin{aligned} V_T &= h(X_T) \quad \text{and} \\ V_t &= \max \left(h(X_t), E_t^{\mathcal{Q}} \left[\frac{B_t}{B_{t+1}} V_{t+1}(X_{t+1}) \right] \right). \end{aligned}$$

The price of the option is then given by $V_0(X_0)$ where X_0 is the initial state of the economy. As an alternative to value iteration we could use *Q-value iteration*. If the Q-value function is defined to be the value of the option conditional on it not being exercised today (i.e. the continuation value of the option) then we also have

$$Q_t(X_t) = E_t^{\mathcal{Q}} \left[\frac{B_t}{B_{t+1}} V_{t+1}(X_{t+1}) \right].$$

The value of the option at time $t + 1$ is then

$$V_{t+1}(X_{t+1}) = \max(h(X_{t+1}), Q_{t+1}(X_{t+1}))$$

so that we can also write

$$Q_t(X_t) = E_t^{\mathcal{Q}} \left[\frac{B_t}{B_{t+1}} \max(h(X_{t+1}), Q_{t+1}(X_{t+1})) \right]. \tag{2}$$

Equation (2) clearly gives a natural analog to value iteration, namely Q-value iteration. As stated earlier, if n is large so X_t is high dimensional, then both value iteration and Q-value iteration are not feasible in practice. However, we could perform an *approximate* and efficient version of Q-value iteration, and this is precisely what the ADP algorithms of Longstaff and Schwartz (2001) and Tsitsiklis and Van Roy (2001) do. We now describe their main contribution, omitting some of the more specific details that can nevertheless have a significant impact on performance.

The first step is to choose a set of *basis functions*, $\phi_1(X), \dots, \phi_m(X)$. These basis functions define the *linear architecture* that will be used to approximate the Q-value functions. In particular, we will approximate $Q_t(X_t)$ with

$$\tilde{Q}_t(X_t) = r_t^1 \phi_1(X) + \dots + r_t^m \phi_m(X)$$

where $r_t := (r_t^1, \dots, r_t^m)$ is a vector of time t parameters that is determined by the algorithm which proceeds as follows:

Approximate Q-Value Iteration

```

generate  $N$  paths of state vector,  $X$ 
set  $\tilde{Q}_T(X_T^i) = 0$  for all  $i = 1$  to  $N$ 
for  $t = T - 1$  downto 1
    Estimate  $r_t = (r_t^1, \dots, r_t^m)$ 
    set  $\tilde{Q}_t(X_t^i) = \sum_k r_t^k \phi_k(X_t^i)$  for all  $i$ 
end for
set  $\tilde{V}_0(X_0) = \max(h(X_0), \tilde{Q}_0(X_0))$ 

```

Two steps require further explanation. First, we estimate r_t by regressing $\alpha \max(h(X_{t+1}), \tilde{Q}(X_{t+1}))$ on $(\phi_1(X_t), \dots, \phi_m(X_t))$ where $\alpha = B_t/B_{t+1}$ is the discount factor. We have N observations for this regression and N is usually taken to be somewhere between 10,000 and 50,000. Second, since all N paths have the same starting point, X_0 , we can estimate $\tilde{Q}_0(X_0)$ by averaging and discounting $\tilde{Q}_1(\cdot)$ evaluated at the N successor points of X_0 .

Obviously many more details are required to fully specify the algorithm. In particular, parameter values and basis functions need to be chosen and specific implementation details can vary. In practice, it is quite common for an alternative estimate, \underline{V}_0 , of V_0 to be obtained by simulating the exercise strategy that is defined implicitly by the sequence of Q-value function approximations. That is, we define $\tilde{\tau} = \min\{t \in \mathcal{T} : \tilde{Q}_t \leq h_t\}$ and

$$\underline{V}_0 = E_0^Q \left[\frac{h_{\tilde{\tau}}}{B_{\tilde{\tau}}} \right].$$

\underline{V}_0 is then an unbiased lower bound on the true value of the option as it is the price that corresponds to a feasible adapted exercise strategy.

These algorithms have performed surprisingly well on realistic high-dimensional problems (see Longstaff and Schwartz 2001 for numerical examples) and there has also been considerable theoretical work (e.g. Tsitsiklis and Van Roy 2001) explaining why this is so. The quality of \underline{V}_0 , for example, can be explained in part by noting that exercise errors are never made as long as $Q_t(\cdot)$ and $\tilde{Q}_t(\cdot)$ lie on the *same* side of the optimal exercise boundary. This means in particular, that it is possible to have large errors in $\tilde{Q}_t(\cdot)$ that do not impact the quality of \underline{V}_0 .

Clearly, simulation plays an important role in these ADP algorithms as it is required to generate the N sample paths of X and to estimate \underline{V}_0 . There are also opportunities for simulation techniques to significantly improve the efficiency of these ADP algorithms as well as the dual-based methods of Section 2.2 that can be used to evaluate ADP solutions.

2.2 Duality Theory for American Options

While ADP methods have been very successful, an important weakness is their inability to determine how far the ADP solution is from optimality in any given problem. Haugh and Kogan (2001) and Rogers (2002) independently developed dual-based methods for evaluating any approximate solution by using it to construct an upper bound on the true value function. (As we saw in Section 2.1, a lower bound is easy to compute. We also remark that Broadie and Glasserman (1997) were the first to demonstrate that tight lower and upper bounds could be constructed using simulation techniques. Their method, however, does not work with arbitrary approximations to the value function and is not as efficient as the dual-ADP techniques.) We now describe these dual based methods.

For an arbitrary adapted supermartingale, π_t , the value of an American option, V_0 , satisfies

$$\begin{aligned} V_0 &= \sup_{\tau \in \mathcal{T}} \mathbb{E}_0^{\mathcal{Q}} \left[\frac{h_\tau}{B_\tau} \right] = \sup_{\tau \in \mathcal{T}} \mathbb{E}_0^{\mathcal{Q}} \left[\frac{h_\tau}{B_\tau} - \pi_\tau + \pi_\tau \right] \\ &\leq \sup_{\tau \in \mathcal{T}} \mathbb{E}_0^{\mathcal{Q}} \left[\frac{h_\tau}{B_\tau} - \pi_\tau \right] + \pi_0 \\ &\leq \mathbb{E}_0^{\mathcal{Q}} \left[\max_{t \in \mathcal{T}} \left(\frac{h_t}{B_t} - \pi_t \right) \right] + \pi_0 \end{aligned} \quad (3)$$

where the first inequality follows from the optional sampling theorem for supermartingales. Taking the infimum over all supermartingales, π_t , on the right hand side of (3) implies

$$V_0 \leq U_0 := \inf_{\pi} \mathbb{E}_0^{\mathcal{Q}} \left[\max_{t \in \mathcal{T}} \left(\frac{h_t}{B_t} - \pi_t \right) \right] + \pi_0 \quad (4)$$

On the other hand, it is known (see e.g. Duffie 1996) that the process V_t/B_t is itself a supermartingale, which implies

$$U_0 \leq \mathbb{E}_0^{\mathcal{Q}} \left[\max_{t \in \mathcal{T}} (h_t/B_t - V_t/B_t) \right] + V_0.$$

Since $V_t \geq h_t$ for all t , we conclude that $U_0 \leq V_0$. Therefore, $V_0 = U_0$, and equality is attained when $\pi_t = V_t/B_t$.

This shows that an upper bound on the price of the American option can be constructed simply by evaluating the right-hand-side of (3) for a given supermartingale, π_t . In particular, if such a supermartingale satisfies $\pi_t \geq h_t/B_t$, the option price V_0 is bounded above by π_0 .

When the supermartingale π_t in (3) coincides with the discounted option value process, V_t/B_t , the upper bound on the right-hand-side of (3) equals the true price of the American option. This suggests that a tight upper bound can be obtained by using an accurate approximation, \tilde{V}_t , to define π_t . One possibility (see Haugh and Kogan 2001, and Andersen and Broadie 2001 for further comments related

to the choice of π_t) is to define π_t as a martingale:

$$\pi_0 = \tilde{V}_0 \tag{5}$$

$$\pi_{t+1} = \pi_t + \frac{\tilde{V}_{t+1}}{B_{t+1}} - \frac{\tilde{V}_t}{B_t} - \mathbb{E}_t \left[\frac{\tilde{V}_{t+1}}{B_{t+1}} - \frac{\tilde{V}_t}{B_t} \right]. \tag{6}$$

Let \bar{V}_0 denote the upper bound we get from (3) corresponding to our choice of supermartingale in (5) and (6). Then it is easy to see that the upper bound is explicitly given by

$$\bar{V}_0 = \tilde{V}_0 + \mathbb{E}_0^{\mathcal{Q}} \left[\max_{t \in \mathcal{T}} \left(\frac{h_t}{B_t} - \frac{\tilde{V}_t}{B_t} + \sum_{j=1}^t \mathbb{E}_{j-1}^{\mathcal{Q}} \left[\frac{\tilde{V}_j}{B_j} - \frac{\tilde{V}_{j-1}}{B_{j-1}} \right] \right) \right]. \tag{7}$$

As may be seen from (7), obtaining an accurate estimate of \bar{V}_0 is computationally demanding. First, a number of sample paths must be simulated to estimate the outermost expectation on the right-hand-side of (7). While this number need not be large in practice, we also need to accurately estimate a conditional expectation at each time period along each simulated path. This requires some effort and clearly variance reduction methods would be useful in this context.

Variations and extensions of these algorithms have also been developed recently and are a subject of ongoing research. Andersen and Broadie (2001), for example, construct upper bounds by using an approximation to the optimal exercise frontier instead of an approximation to the Q-value function, while Meinshausen and Hambly (2003) use similar ideas to price options that may be exercised multiple times.

3 PORTFOLIO OPTIMIZATION

Motivated by the success of ADP methods for pricing American options, Brandt et al (2001) apply similar ideas to approximately solve a class of high-dimensional portfolio optimization problems. In particular, they simulate a large number of sample paths of the underlying state variables and then working backwards in time, they use *cross path regressions* (as we described in the approximate Q-value iteration algorithm) to efficiently compute an approximately optimal strategy. Propagation of errors is largely avoided, and though the price for this is an algorithm that is quadratic in the number of time periods, their methodology can comfortably handle problems with a large number of time periods. Their specific algorithm does not handle portfolio constraints and certain other complicating features, but it should be possible to tackle these extensions using the ADP methods that they and others have developed.

As was the case with ADP solutions to optimal stopping problems, a principal weakness of ADP solutions to portfolio optimization problems is the difficulty in determining how far a given solution to a given problem is from optimality. This issue has motivated in part the research of Haugh, Kogan and Wang (2003) (hereafter HKW) who use portfolio duality theory to evaluate the quality of suboptimal solutions to portfolio optimization problems by constructing lower and upper bounds on the optimal

value function. These bounds are evaluated by simulating the stochastic differential equations (see Kloeden and Platen 1992) that describe the evolution of the state variables in the model in question. In section 3.1 we describe the particular portfolio duality theory that was used in HKW and that was developed by Xu (1990), Shreve and Xu (1992a, 1992b), Karatzas, Lehocky, Shreve and Xu (1991), and Cvitanic and Karatzas (1992).

Before doing so, we remark on the role that simulation has to play in applying this theory in practice. First, since the portfolio optimization problems in question are too difficult to solve either analytically or numerically, approximate solution techniques are necessary. To date, ADP methods appear to be the most promising and as we have seen, simulation plays an important role when applying these methods. Second, once we have an approximate solution we would like to evaluate it by using the solution itself to construct lower and upper bounds on the true value function. Again, we can only do this by simulating paths of the relevant state variables as will be discussed at the end of Section 3.1.

We also remark that duality theory of Section 3.1 applies mainly to problems in continuous time. ADP techniques, on the other hand, are generally more suited to a discrete time framework, This inconsistency can easily be overcome by extrapolating discrete-time ADP solutions to construct continuous-time solutions.

3.1 Dual Methods for Portfolio Optimization

We assume that the financial market has N stocks whose time t prices are given by the N -vector, P_t , and a cash account that earns interest at the instantaneously risk-free rate, r_t . We let the M -vector X_t denote the time t value of the state variables in the market. The dynamics of these variables is governed by the following system of stochastic differential equations (SDE's):

$$r_t = r(X_t) \tag{8a}$$

$$dP_t = P_t [\mu_P(X_t) dt + \Sigma_P(X_t) dB_t] \tag{8b}$$

$$dX_t = \mu_X(X_t) dt + \Sigma_X(X_t) dB_t \tag{8c}$$

where $B_t = (B_{1t}, \dots, B_{Nt})$ is a vector of N independent Brownian motions, μ_P and μ_X are N - and M -dimensional drift vectors, and Σ_P and Σ_X are diffusion matrices of dimension N by N and M by N , respectively. Without loss of generality, we assume that $M < N$ and that the last M rows of (8b) coincide with (8c). We define $\eta_t = \Sigma_{P_t}^{-1}(\mu_{P_t} - r_t)$ so that in a market without portfolio constraints, η_t corresponds to the vector *market price of risk*.

The time t portfolio weights are denoted by $\theta_t = (\theta_{1t}, \dots, \theta_{Nt})$ and the interpretation is that θ_{it} is the fraction of wealth, W_t , that is invested in the i^{th} stock at time t . (This implies $1 - \sum \theta_{it}$ is invested in the cash account at time t .) The wealth dynamics are then given by (e.g. Duffie 1996)

$$dW_t = W_t \left\{ \left[r_t + \theta_t^\top (\mu_{P_t} - r_t) \right] dt + \theta_t^\top \Sigma_{P_t} dB_t \right\} \tag{9}$$

and for ease of exposition, we are now dropping the dependence of terms on X_t .

To describe the constraints that a portfolio strategy must satisfy, we let \mathbf{K} be a closed convex set in

R^N that contains the 0 vector. We assume that θ_t must satisfy

$$\theta_t \in \mathbf{K} \quad \text{for all } t \in [0, T]. \quad (10)$$

For example, if short sales are not allowed, then $\mathbf{K} = \{\theta : \theta \geq 0\}$. If in addition borrowing is also not allowed, then $\mathbf{K} = \{\theta : \theta \geq 0, \mathbf{1}^\top \theta \leq 1\}$.

The portfolio optimization problem is to maximize expected utility of terminal wealth, $E[U(W_T)]$. In particular, we must solve for

$$V_0 \equiv \sup_{\{\theta_t\}} E_0 [U(W_T)] \quad (\mathcal{P})$$

subject to (8), (9) and (10)

where V_0 denotes the value function at $t = 0$ and where the initial wealth is assumed to be W_0 .

Because the number of Brownian motions, N , is equal to the number of stocks in the financial market described by (8), it can be shown that the market would be a *dynamically complete* market if there were no portfolio constraints. Dynamic completeness would imply the existence of a unique market-price-of-risk process, η_t , or equivalently, a unique *state-price-density* (SPD) process, π_t . $\Pi_t(\omega)$ may be interpreted as the price per-unit-probability of \$1 at time t in the event ω occurs. (See Duffie 1996 for further details.)

It so happens that a portfolio optimization problem in complete markets is particularly easy to solve using martingale methods as the problem can essentially be decoupled (Cox and Huang 1991, and Karatzas, Lehocky and Shreve 1997). First the optimal wealth, W_T^* , is chosen in such a way that the budget constraint,

$$E_0 [\Pi_T W_T^*] = W_0$$

is satisfied. The second step is to solve for the portfolio strategy, θ^* , that attains W_T^* . This decoupling is not possible in incomplete markets since it is not the case that every random terminal wealth, W_T , is attainable using a self-financing trading strategy. In problem (\mathcal{P}) above, the imposition of portfolio constraints implies that we do not have a complete financial market and so the simple martingale decoupling approach cannot be used. In general, we are then left with a problem that cannot be solved explicitly. Moreover, it might be very difficult to solve the problem numerically and it will be impossible to do so if the problem is also high-dimensional.

In recent years, however, the martingale approach for complete markets has been generalized to incomplete markets by a number of researchers using stochastic duality theory to allow for portfolio constraints and non-spanned risks. Research in this direction includes He and Pearson (1991), Karatzas, Lehoczky, Shreve, and Xu (1991), Cvitanic and Karatzas (1992), Cuoco (1997), and Cuoco and Liu (2000), to name a few. Rogers (2003) provides a synthesis of many of the results to date. Explicit solutions to these problems are rare and notable exceptions are problems with logarithmic preferences, or problems with constant relative risk aversion (CRRA) preferences and a deterministic investment opportunity set (see, for example, Karatzas and Shreve 1997, Section 6.6).

All is not lost, however, for even though explicit solutions are rare, it might still be possible to use this duality theory in practice. In particular, HKW show how the duality theory of Cvitanic and Karatzas (1992) and others may be used to evaluate approximate solutions to difficult portfolio optimization problems. Their methods, which apply to problems in a multidimensional diffusion setting, should be of value when exact solutions are not available.

Starting with the portfolio choice problem (\mathcal{P}), we can define a fictitious problem ($\mathcal{P}^{(\nu)}$), based on a fictitious financial market and *without* the portfolio constraints. First we define the *support function* of \mathbf{K} , $\delta(\cdot) : \mathbf{R}^N \rightarrow \mathbf{R} \cup \infty$, by

$$\delta(\nu) \equiv \sup_{x \in \mathbf{K}} (-\nu^\top x). \quad (11)$$

The effective domain of the support function is given by

$$\tilde{\mathbf{K}} \equiv \{\nu : \delta(\nu) < \infty\}. \quad (12)$$

For suitably well-behaved processes, ν_t , that satisfy $\delta(\nu) < \infty$ almost surely for all t , we define a fictitious market $M^{(\nu)}$, in which the N stocks and the cash account are traded without constraints. The diffusion matrix of stock returns in $M^{(\nu)}$ is the same as in the original market. However, the risk-free rate and the vector of expected stock returns are different. In particular, the risk-free rate process and the price drift vector in the fictitious market are defined respectively by

$$\begin{aligned} r_t^{(\nu)} &= r_t + \delta(\nu_t) \\ \mu_{P_t}^{(\nu)} &= \mu_{P_t} + \delta(\nu_t) + \nu_t \end{aligned}$$

where $\delta(\nu)$ is the support function defined in (11). The dynamic portfolio optimization problem in the fictitious market, $M^{(\nu)}$, *without portfolio constraints* is a complete markets problem and so there exists a unique SPD process, $\pi_t^{(\nu)}$. As stated above, the decoupling approach can be used to solve this problem. In particular, it may be formulated as a static problem as follows:

$$V_0^{(\nu)} \equiv \sup_{\{W_T\}} \mathbb{E}_0 [U(W_T)] \quad (\mathcal{P}^{(\nu)})$$

$$\text{subject to } \mathbb{E}_0 \left[\pi_T^{(\nu)} W_T \right] = W_0.$$

(Note that in problem $\mathcal{P}^{(\nu)}$ we focus on finding the optimal terminal wealth, W_T , and do not need to worry about the optimal strategy, θ_t .) Due to its static nature, the problem ($\mathcal{P}^{(\nu)}$) is easy to solve. For example, when the utility function is of the form $U(W) = W^{1-\gamma}/(1-\gamma)$, the corresponding value function in the fictitious market is given explicitly by

$$V_0^{(\nu)} = \frac{W_0^{1-\gamma}}{1-\gamma} \mathbb{E}_0 \left[\left[\pi_T^{(\nu)} \right]^{\frac{\gamma-1}{\gamma}} \right]^\gamma. \quad (14)$$

It is easy to see that for any choice of ν , the value function in (14) gives an upper bound for the optimal value function, V_0 , of the original problem. In the fictitious market, the wealth dynamics of the portfolio are given by

$$dW_t^{(\nu)} = W_t^{(\nu)} \left[\left(r_t^{(\nu)} + \theta_t^\top (\mu_{P_t}^{(\nu)} - r_t^{(\nu)}) \right) dt + \theta_t^\top \Sigma_{P_t} dB_t \right]$$

so that

$$\frac{dW_t^{(\nu)}}{W_t^{(\nu)}} - \frac{dW_t}{W_t} = \left[\delta(\nu_t) + \theta_t^\top \nu_t \right] dt.$$

If $\theta_t \in \mathbf{K}$ then (11) implies that the last expression is non-negative. Therefore $W_t^{(\nu)} \geq W_t \forall t \in [0, T]$ and since $U(\cdot)$ is assumed to be an increasing function, we have $V_0^{(\nu)} \geq V_0$. Clearly we also have

$$\inf_{(\nu)} V_0^{(\nu)} \geq V_0. \quad (15)$$

Cvitanic and Karatzas (1992) and other researchers (e.g. Schroder and Skiadis 2003) have shown that in many circumstances there is no duality gap. That is, there exists ν^* such that $V_0^{(\nu^*)} = V_0$. HKW show then that

$$\eta_t^{(\nu^*)} = -W_t \frac{\partial^2 V_t / \partial W_t^2}{\partial V_t / \partial W_t} \Sigma_{P_t}^\top \theta_t^* - \left(\frac{\partial V_t}{\partial W_t} \right)^{-1} \Sigma_{X_t}^\top \left(\frac{\partial^2 V_t}{\partial W_t \partial X_t} \right)$$

where θ_t^* denotes the optimal portfolio policy for the original problem and $\eta_t^{(\nu^*)}$ is the *market-price-of-risk* process in the optimal fictitious market. We remark that the important feature of (16) is that it provides a link between the primal and dual optimal solutions. Since finding optimal solutions to these problems is generally not possible, we have to make do with finding suboptimal solutions, \tilde{V}_t and $\tilde{\theta}_t$. We can evaluate such a suboptimal solution by using it to construct a particular fictitious market. In particular, we can replace V_t and θ_t^* in (16) with \tilde{V}_t and $\tilde{\theta}_t$ to obtain

$$\eta_t^{(\tilde{\nu})} = -W_t \left(\frac{\partial_{WW} \tilde{V}_t}{\partial_W \tilde{V}_t} \right) \Sigma_{P_t}^\top \tilde{\theta}_t - \left(\partial_W \tilde{V}_t \right)^{-1} \Sigma_{X_t}^\top \left(\partial_{WX} \tilde{V}_t \right).$$

We can then use $\eta_t^{(\tilde{\nu})}$ to define $\tilde{\nu}_t$. If these processes are suitably well behaved (e.g. $\delta(\tilde{\nu}_t) < \infty$ almost surely for all t) they can be used to define an upper bound on the true value function, V_0 . This upper bound is given by equation (14) and it is shown in HKW how a good approximate primal solution, $\tilde{\theta}_t$, leads to a good upper bound on V_0 . (Computing a lower bound is generally easy: simply simulate the policy implicitly defined by the initial approximation.)

We also mention that in the special but important case of a CRRA (constant relative risk aversion) utility function, the expression for $\eta_t^{(\nu^*)}$ can be shown to simplify considerably. In particular, the first term in (16) is then a function of θ_t^* only, while approximations to the value function and only its first derivatives with respect to the state variables, X_t , are needed to evaluate the second term. This simplifies numerical implementation significantly, but of course approximating functions and their partial

derivatives is still very challenging in practice.

Simulation: The expectation in (14) cannot be evaluated explicitly and so it has to be estimated by simulating the underlying SDE's. This is a computationally intensive task, particularly when \tilde{v}_t cannot be guaranteed in advance to be well-behaved. In such circumstances it is necessary to solve a quadratic optimization problem at each discretization point on each simulated path in order to convert $\eta_t^{(\tilde{v})}$ and \tilde{v}_t into well-behaved versions that can then be used to construct an upper bound on V_0 . (See HKW for further details.)

Besides the actual ADP implementation that constructs the initial approximate solution, simulation is also often necessary to approximate the value function and its partial derivatives in (16). This occurs when we wish to evaluate a given portfolio policy, $\tilde{\theta}_t$, but do not know the corresponding value function, \tilde{V}_t . In such circumstances, it seems that it is necessary to simulate the policy, $\tilde{\theta}_t$, in order to approximate the required functions. Once again, this is computationally demanding and seeking efficient simulation techniques for all of these tasks will be an important challenge as we seek to solve ever more complex problems.

4 CONCLUSIONS

Simulation has an important role to play in the optimal control problems that are found in financial engineering. These control problems include optimal stopping and portfolio optimization problems which in many circumstance cannot be solved exactly. When this is the case, approximate solution techniques are required and the most successful techniques to date are probably the approximate dynamic programming (ADP) techniques. A particular weakness with the ADP methodology, however, is that it is difficult to determine how far a particular solution to a particular problem is from optimality. Dual-based methods have recently proved useful for evaluating approximate solutions by enabling the computation of lower *and* upper bounds on the true value function.

Simulation is a necessary and important tool for constructing the initial solution as well as evaluating it by computing lower and upper bounds. Since all of these methods are computationally intensive, it is expected that more sophisticated simulation techniques will have a greater role to play in future research. Moreover, in the context of portfolio optimization, there are many different 'formulations' of the duality theory (see Rogers 2003), and it is expected that many of these formulations can be used in a computational framework just as the dual formulation of Cvitanic and Karatzas (1992) was used in HKW (2003). This is a topic of ongoing research and simulation techniques will certainly be an important tool in furthering this agenda.

REFERENCES

- Andersen, L. and M. Broadie. 2001. A Primal-Dual Simulation Algorithm for Pricing Multi-Dimensional American Options. Working Paper, Columbia University.
- Bertsekas, D., and J. Tsitsiklis. 1996. *Neuro-Dynamic Programming*. Belmont, Massachusetts: Athena Scientific .

- Black, F., and M. Scholes. 1973. The Pricing of Options and Corporate Liabilities. *Journal of Political Economy* 81:637-654.
- Boyle, P., M. Broadie, and P. Glasserman. 1997. Monte Carlo Methods for Security Pricing. *Journal of Economic Dynamics and Control* 21:1267-1321.
- Brandt, M.W., A. Goyal, P. Santa-Clara, and J.R. Stroud. 2001. A Simulation Approach to Dynamic Portfolio Choice with an Application to Learning in Discrete Time. Working paper, Duke University.
- Broadie, M., and P. Glasserman. 1997. A Stochastic Mesh Method for Pricing High-Dimensional American Options. Working Paper, Columbia Business School, Columbia University, New York.
- Cox, J., and C.-F. Huang. 1989. Optimal Consumption and Portfolio Policies When Asset Prices Follow a Diffusion Process. *Journal of Economic Theory* 49:33-83.
- Cuoco, D. 1997. Optimal Consumption and Equilibrium Prices with Portfolio Constraints and Stochastic Income. *Journal of Economic Theory* 72:33-73.
- Cuoco, D. and H. Liu. 2000. A Martingale Characterization of Consumption Choices and Hedging Costs with Margin Requirements. *Mathematical Finance* 10:355-385.
- Cvitanić, J., and I. Karatzas. 1992. Convex Duality in Constrained Portfolio Optimization. *Annals of Applied Probability* 2:767-818.
- Duffie, D. 1996. *Dynamic Asset Pricing Theory*. Princeton, New Jersey: Princeton University Press.
- Haugh, M.B., and L. Kogan. 2001. Pricing American Options: A Duality Approach. Forthcoming in *Operations Research*. <<http://www.columbia.edu/~mh2078/Research.html>>.
- Haugh, M.B., L. Kogan and J. Wang. 2003. Portfolio Evaluation: A Duality Approach. Working paper, Department of IE & OR, Columbia University. <<http://www.columbia.edu/~mh2078/Research.html>>.
- He, H., and N. Pearson. 1991. Consumption and Portfolio Policies with Incomplete Markets and Short-sale Constraints: The Infinite Dimensional Case. *Journal of Economic Theory* 54 (2):259-304.
- Karatzas, I., J.P. Lehocky and S.E. Shreve. 1987. Optimal Portfolio and Consumption Decisions for a "Small Investor" on a Finite Horizon. *SIAM J. Control Optimization* 25:1557-1586.
- Karatzas, I., J.P. Lehocky, S.E. Shreve and G.L. Xu. 1991. Martingale and Duality Methods for Utility Maximization in an Incomplete Market. *SIAM J. Control Optimization* 29:702-730.
- Karatzas, I., and S.E. Shreve. 1997. *Methods of Mathematical Finance*. New York: Springer-Verlag.
- Kloeden, P., and E. Platen. 1992. *Numerical Solution of Stochastic Differential Equations*. Berlin: Springer-Verlag.
- Liu, J. 1998. Dynamic Portfolio Choice and Risk Aversion. Working paper, Stanford University, Palo Alto.
- Longstaff, F., and E. Schwartz. 2001. Valuing American Options by Simulation: A Simple Least-Squares Approach. *Review of Financial Studies* 14:113-147.

- Meinshausen, N., and B.M. Hambly. 2003. Monte Carlo Methods for the Valuation of Multiple Exercise Options. Working paper, Oxford University.
- Merton, R.C. 1990. *Continuous-Time Finance*. New York: Basil Blackwell.
- Rogers, L.C.G. 2002. Monte-Carlo Valuation of American Options. *Mathematical Finance*. 12(3):271-286.
- Rogers, L.C.G. 2003. Duality in Constrained Optimal investment and Consumption Problems: A Synthesis. Working paper, Statistical Laboratory, Cambridge University. <<http://www.statslab.cam.ac.uk/~chris/>>.
- Schroder, M. and C. Skiadas. 2003 Optimal Lifetime Consumption-Portfolio Strategies under Trading Constraints and Generalized Recursive Preferences. Working Paper No. 285. Kellogg School of Management, Northwestern University, Evanston.
- Shreve, S.E., and G.L. Xu. 1992a. A Duality Method for Optimal Consumption and Investment under Short-Selling Prohibition, Part I: General Market Coefficients. *Annals of Applied Probability* 2:8-112.
- Shreve, S.E., and G.L. Xu. 1992b. A Duality Method for Optimal Consumption and Investment under Short-Selling Prohibition, Part I: Constant Market Coefficients. *Annals of Applied Probability* 2:314-328.
- Staum, J. 2002. Simulation in Financial Engineering. In *Proceedings of the 2002 Winter Simulation Conference*, ed. E. Yücesan, C.-H. Chen, J.L. Snowdon, and J.M. Charnes, 1481-1492.
- Tsitsiklis, J., and B. Van Roy. 2001. Regression Methods for Pricing Complex American-Style Options. *IEEE Transactions on Neural Networks* 12(4):694-703.
- Xu, G.L. 1990. A Duality Method for Optimal Consumption and Investment under Short-Selling Prohibition. Ph.D. Dissertation, Carnegie Mellon University.