Evaluating Portfolio Policies: A Duality Approach

Martin B. Haugh

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

Leonid Kogan

Sloan School of Management, MIT, Massachusetts, MA 02142, `lkogan@mit.edu`

Jiang Wang

Sloan School of Management, MIT, Massachusetts, MA 02142, CCFR and NBER, `wangj@mit.edu`

The performance of a given portfolio policy can in principle be evaluated by comparing its expected utility with that of the optimal policy. Unfortunately, the optimal policy is usually not computable in which case a direct comparison is impossible. In this paper we solve this problem by using the given portfolio policy to construct an upper bound on the unknown maximum expected utility. This construction is based on a dual formulation of the portfolio optimization problem. When the upper bound is close to the expected utility achieved by the given portfolio policy, the potential utility loss of this policy is guaranteed to be small. Our algorithm can be used to evaluate portfolio policies in models with incomplete markets and position constraints. We illustrate our methodology by analyzing the static and myopic policies in markets with return predictability and constraints on short sales and borrowing.

1 Introduction

Optimal portfolio choice is one of the central subjects in finance (see e.g., Merton, 1990). Yet analytical solutions are known only in a few special cases, under restrictive assumptions on the market structure and/or the investor's utility function. Most problems of practical interest cannot be solved in closed form and must be analyzed using approximations. But how costly are the approximation errors of a candidate portfolio policy to the investor? To answer this question, one must compare the expected utility under the candidate policy with the expected utility under the optimal policy. Such a direct comparison is impossible when the optimal policy cannot be computed explicitly. In this paper, we construct an upper bound on the expected utility under the optimal policy based on the dual formulation of the portfolio optimization problem developed by Cvitanic and Karatzas (1992). When the upper bound is close to the expected utility achieved by a candidate portfolio policy, the utility cost of choosing such a suboptimal policy is guaranteed to be small. Our method applies to models with incomplete markets and position constraints.

The most general class of tractable problems are those where the financial market is dynamically complete. Market completeness allows one to use the martingale techniques to restate the portfolio choice problem in an equivalent static form. This was first shown by Cox and Huang (1989), Karatzas, Lehoczky, Sethi, and Pliska (1986) and Shreve (1986). The static problem often leads to analytical solutions and is also amenable to simulation techniques (see e.g., Detemple, Garcia, and Rindisbacher 2000). When the financial market is incomplete due to position constraints or non-spanned risks, explicit solutions have been obtained for only a few special cases. For example, Kim and Omberg (1996) solve the portfolio choice problem assuming an affine non-spanned process for the market price of risk and no portfolio constraints (with power utility function). Liu (2005) extends their results to multi-variate problems.

The martingale approach has been generalized by a number of researchers using stochastic duality theory to allow for portfolio constraints and non-spanned risks. They include Cuoco (1997), Cvitanic and Karatzas (1992), He and Pearson (1991), and Karatzas, Lehoczky, Shreve, and Xu (1991), among others. However, explicit solutions are rare. Notable exceptions are problems with logarithmic preferences, or problems with constant relative risk aversion (CRRA) preferences and deterministic investment opportunity set (see, for example, Karatzas and Shreve 1998, Section 6.6). For most portfolio choice problems, explicit solutions are not available and one must rely on numerical approximations. Various types of approximate methods have been used in the literature. They include the log-linear analytical approximations of Campbell and Viceira (1999, 2002) (see also Campbell, Chan and Viceira 2002 and Chacko and Viceira 2002), finite-difference PDE solution methods (e.g., Brennan 1998, Brennan, Schwartz, and Lagnado 1997, Brennan and Xia 2001, and Xia 2001), Markov-chain approximations (e.g., Munk, 1999), numerical dynamic programming based on state-space discretization for discrete-time problems (e.g., Ait-Sahalia and Brandt 2001, Balduzzi and Lynch 1999, Brandt 1999, Barberis 2000, Dammon, Spatt, and Zhang 2000, and Lynch 2001) and the approximate dynamic programming approach of Brandt, Goyal, and Santa-Clara (2001). Other techniques for solving dynamic optimization problems that are applicable to portfolio choice problems are described in Judd (1996) and Rust (1996).

An important limitation of the approximate solution methods is that it is usually difficult to evaluate the accuracy of the obtained approximations. Judd (1996) points out that there are few theoretical results on the convergence properties of the numerical algorithms used in economic applications and even when such results are available, asymptotic convergence does not guarantee

the accuracy of any given approximation. Moreover, some of the existing approaches are based on a single approximation, rather than a sequence of successive approximations (e.g., the log-normal approximation method used by Campbell and Viceira, 1999), and for these the notion of asymptotic convergence is not even applicable. Den Haan and Marcet (1994) and Judd (1996) suggest procedures for testing how accurately the first-order optimality conditions are satisfied. Such tests, however, provide little information about the economic impact of the approximation errors.

The method we propose allows us to evaluate the quality of approximations for portfolio choice problems by providing an upper bound on the utility loss associated with an approximate solution. If a sequence of progressively refined approximate portfolio policies were generated using an iterative numerical algorithm then our method could be used to determine at what point the quality of the approximation becomes acceptable. At that point the algorithm could then be terminated. Since our approach does not depend on the particular method used to construct the approximate policy and value function, it can be applied to any approximate solution method. Moreover, in principle our methodology could be used to guide the numerical algorithm as well. In this paper, however, we will focus on evaluating the quality of a fixed portfolio policy and not attempt to construct and evaluate a *sequence* of approximations.

Our approach is based on the duality formulation of the portfolio choice problem proposed by Cvitanic and Karatzas (1992). For a given portfolio choice problem, which is subject to position constraints, we can define a fictitious, *unconstrained* problem with modified price processes such that the maximum expected utility achieved in the fictitious problem is at least as high as in the original problem. Under certain assumptions, the constructed upper bound can be shown to be tight, i.e., to coincide with the maximum expected utility for the original problem. Since the fictitious problem is unconstrained, we can solve it explicitly using the well-known martingale techniques of Cox and Huang (1989). The solution gives an upper bound on the maximum expected utility (the value function) of the original problem. (Haugh and Kogan, 2003, develop a similar algorithm for the problem of pricing American options, which is an optimal stopping problem. As in this paper, the upper bound on the value function is constructed using a dual formulation of the original problem.)

In order to evaluate an approximation for the optimal policy and value function, we can use this upper bound when the value function itself is not available. The goal is to make the upper bound tight by properly choosing the fictitious problem, in particular the fictitious market with the modified price processes.

Our key insight is to perform such a modification using the information contained in the approximate solution to the problem. In particular, one must first specify the approximate portfolio policy and the value function. It would be natural to compute the value function implied by the approximate portfolio policy, but it is not necessary to do so. In our numerical examples below, we specify an approximate value function directly in a particularly simple way and use it together with different portfolio policies to establish bounds on the true value function. The approximation to the value function is then used to define the candidate return processes for all the assets in the fictitious financial market, as implied by the optimality conditions of the original portfolio choice problem. The implied return processes in general do not qualify to be the return processes in the fictitious problem. Our next step is to choose the qualified fictitious return processes that are "closest" to the candidate processes. We then solve the unconstrained optimization problem in the

fictitious market. The value function of this fictitious problem gives an upper bound on the value function of the original problem.

We illustrate our method with a few simple numerical examples. These are not designed to be fully realistic, but rather serve to demonstrate the potential of our algorithm. Since it is not the purpose of this paper to develop new schemes for computing approximate portfolio policies, to simplify implementation, we adopt very simple, analytical approximations to the portfolio policies. In particular, we evaluate the quality of static and myopic portfolio policies in a market with return predictability under constraints on short sales and borrowing. A static portfolio policy is obtained by ignoring time-variation in the investment opportunity set and fixing the instantaneous moments of stock returns at their long-run average values. To obtain a myopic policy, we assume at any point in time that the investment opportunity will remain constant thereafter. By comparing the expected utility under the static strategy with our upper bound one can measure the economic importance of return predictability. A small gap between the upper bound and the value function under the static policy would indicate that return predictability is not economically important for the problem at hand.

Similar logic can be applied to the myopic policy. A small gap under the myopic policy indicates that the hedging component of the optimal portfolio policy has little impact on the expected utility. The analysis of the myopic policy is particularly relevant since the existing literature analyzing the patterns of predictability in stock returns commonly ignores the hedging component of portfolio policies (see, for example, Johannes, Polson and Stroud 2002, Kandel and Stambaugh 1996, Pastor 2000, Pastor and Stambaugh 2000, and Stambaugh 1999). This is motivated by the fact that for most realistic problems the optimal portfolio policy cannot be computed numerically due to the curse of dimensionality and analytical solutions are only available for highly restrictive settings, which typically rule out portfolio constraints and limit the range of possible specifications for the return processes. Several authors have argued in the context of specific applications that hedging demand is relatively small (e.g., Ait-Sahalia and Brandt 2001, Ang and Bekaert, 2000, Brandt 1999, Chacko and Viceira 2002 and Gomes 2001). Based on this, one may be tempted to assume that for "similar" problems, hedging demand can be safely ignored. However, we know that this cannot be the case in general and for important classes of problems hedging demand is crucial. These classes include long-horizon problems (Barberis 2000, and Campbell and Viceira 1999, 2002) and problems with non-linear dynamics or non-standard preferences (Chan and Kogan 2002). Using our method, one can verify that for a particular problem under consideration ignoring the hedging component of the optimal portfolio policy has little effect on the expected utility.

The rest of the paper is organized as follows. Section 2 formulates the portfolio choice problem. Section 3 reviews the duality theory while Section 4 describes the algorithm for constructing the upper bound on the optimal expected utility. Section 5 illustrates our approach for several types of portfolio choice problems and approximate strategies. We conclude in Section 6.

2   The Model

We now state a portfolio choice problem under incomplete markets and portfolio constraints. We formulate our problem in continuous time and assume that stock prices follow diffusion processes.

The investment opportunity set

There are N stocks and an instantaneously riskfree bond. The vector of stock prices is denoted by $P_t = (P_{1t}, \ldots, P_{Nt})$ and the instantaneously riskfree rate of return on the bond is denoted by r_t. Without loss of generality, we assume that the stocks pay no dividends. The instantaneous moments of asset returns depend on the M-dimensional vector of state variables X_t:

$$r_t = r(X_t) \tag{1a}$$

$$dP_t = P_t \left[\mu_P(X_t) \, dt + \Sigma_P(X_t) \, dB_t \right] \tag{1b}$$

$$dX_t = \mu_X(X_t) \, dt + \Sigma_X(X_t) \, dB_t \tag{1c}$$

where $P_0 = X_0 = 0$, $B_t = (B_{1t}, \ldots, B_{Nt})$ 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. We assume that the diffusion matrix of the stock return process Σ_P is lower-triangular and non-degenerate: $x^\top \Sigma_P \Sigma_P^\top x \geq \epsilon \|x\|^2$ for all x and some $\epsilon > 0$. Then, one can define a process η_t, given by

$$\eta_t = \Sigma_{Pt}^{-1}(\mu_{Pt} - r_t)$$

In a market without portfolio constraints, η_t corresponds to the vector of instantaneous market prices of risk of the N stocks (see, e.g., Duffie 1996, Section 6.G). We adopt a standard assumption that our return generating process is sufficiently well behaved, so that the process η_t is square-integrable:

$$\mathrm{E}_0 \left[\int_0^T \|\eta_t\|^2 \, dt \right] < \infty.$$

Portfolio constraints

A portfolio consists of positions in the N stocks and the riskfree bond. We denote the proportional holdings of risky assets in the total portfolio value by $\theta_t = (\theta_{1t}, \ldots, \theta_{Nt})$. To rule out arbitrage, we require the portfolio policy to satisfy a square integrability condition: $\int_0^T \|\theta_t\|^2 \, dt < \infty$ almost surely. The value of the portfolio changes according to

$$dW_t = W_t \left\{ \left[r_t + \theta_t^\top (\mu_{Pt} - r_t) \right] dt + \theta_t^\top \Sigma_{Pt} \, dB_t \right\}. \tag{2}$$

We assume that the portfolio shares are restricted to lie in a closed convex set $\mathbf{K}$, containing the zero vector:

$$\theta_t \in \mathbf{K}. \tag{3}$$

For example, if short sales are not allowed, then the constraint set takes form $\mathbf{K} = \{\theta : \theta \geq 0\}$. If in addition to prohibiting short sales borrowing is not allowed, then $\mathbf{K} = \{\theta : \theta \geq 0, \ 1^\top \theta \leq 1\}$ where $1^\top = (1, \ldots, 1)$. In our analysis we take the set $\mathbf{K}$ to be constant, but it can be allowed to depend on time and the values of the exogenous state variables.

The objective function

We assume that the portfolio policy is chosen to maximize the expected utility of wealth at the terminal date T, $\mathrm{E}_0[U(W_T)]$. The function $U(W)$ is assumed to be strictly monotone with positive slope, concave and smooth. Moreover, it is assumed to satisfy the Inada conditions at zero and infinity: $\lim_{W \to 0} U'(W) = \infty$ and $\lim_{W \to \infty} U'(W) = 0$. In our numerical examples, we use the utility function with constant relative risk aversion (CRRA) so that $U(W) = W^{1-\gamma}/(1-\gamma)$.

In summary, the portfolio choice problem is to solve for

$$V_0 \equiv \sup_{\{\theta_t\}} \ \mathrm{E}_0\left[U(W_T)\right] \quad \text{subject to} \quad (1), (2) \text{ and } (3) \tag{$\mathcal{P}$}$$

where V_0 denotes the value function at 0.

In the setting we adopt, with continuous-time and diffusion processes for stock prices, the financial market has the property that it can be made dynamically complete by introducing a small number of additional securities. Our analysis primarily relies on this property and thus, can be extended to discrete-time and discrete-state models, using the results in Pliska (1997). However, our approach would not be directly applicable to models which do not possess this property, for instance to models with continuously distributed jumps in stock prices.

3 Review of the Duality Theory

In this section we briefly review the duality approach to the constrained portfolio optimization problem, following Cvitanic and Karatzas (1992) and Schroder and Skiadas (2003). We build on these theoretical results to construct bounds on the performance of portfolio policies in Section 4 below.

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

$$\delta(\nu) \equiv \sup_{x \in \mathbf{K}} (-\nu^\top x). \tag{4}$$

The effective domain of the support function is given by

$$\widetilde{\mathbf{K}} \equiv \{\nu : \ \delta(\nu) < \infty\}. \tag{5}$$

Because the constraint set $\mathbf{K}$ is convex and contains zero, the support function is continuous and bounded from below on its effective domain $\widetilde{\mathbf{K}}$. We then define the set $\mathbf{D}$ of $\mathcal{F}_t$-adapted $\mathbb{R}^N$ valued processes to be

$$\mathbf{D} \equiv \left\{ \nu_t, 0 \leq t \leq T : \ \nu_t \in \widetilde{\mathbf{K}}, \ \mathrm{E}_0\left[\int_0^T \delta(\nu_t)\, dt\right] + \mathrm{E}_0\left[\int_0^T \|\nu_t\|^2\, dt\right] < \infty \right\}. \tag{6}$$

For each process ν in $\mathbf{D}$, we define a fictitious market $M^{(\nu)}$, in which the N stocks and the riskfree bond are traded. The diffusion matrix of stock returns in $M^{(\nu)}$ is the same as in the original market.

However, the riskfree rate and the vector of expected stock returns are different. In particular, the riskfree rate process and the market price of risk in the fictitious market are defined respectively by

$$r_t^{(\nu)} = r_t + \delta(\nu_t) \tag{7a}$$

$$\eta_t^{(\nu)} = \eta_t + \Sigma_{Pt}^{-1} \nu_t \tag{7b}$$

where $\delta(\nu)$ is the support function defined in (4). We assume that $\eta_t^{(\nu)}$ is square-integrable. Following Cox and Huang (1989), the state-price density process $\pi_t^{(\nu)}$ in the fictitious market is given by

$$\pi_t^{(\nu)} = \exp\left(-\int_0^t r_s^{(\nu)}\, ds - \frac{1}{2}\int_0^t \eta_s^{(\nu)\top} \eta_s^{(\nu)}\, ds - \int_0^t \eta_s^{(\nu)\top}\, dB_s\right) \tag{8}$$

and the vector of expected returns is given by

$$\mu_{Pt}^{(\nu)} = r_t^{(\nu)} + \Sigma_{Pt}\, \eta_t^{(\nu)}.$$

The dynamic portfolio choice problem in the fictitious market without position constraints can be equivalently formulated in a static form (e.g., Cox and Huang, 1989, Karatzas and Shreve, 1998, Section 3):

$$V^{(\nu)} \equiv \sup_{\{W_T\}} \mathrm{E}_0\left[U(W_T)\right] \quad \text{subject to} \quad \mathrm{E}_0\left[\pi_T^{(\nu)} W_T\right] \leq W_0. \tag{$\mathcal{P}^{(\nu)}$}$$

Due to its static nature, the problem $(\mathcal{P}^{(\nu)})$ is easy to solve. For example, when the utility function is of the CRRA type with relative risk aversion γ so that $U(W) = W^{1-\gamma}/(1-\gamma)$, the corresponding value function in the fictitious market is given explicitly by

$$V_0^{(\nu)} = \frac{W_0^{1-\gamma}}{1-\gamma}\ \mathrm{E}_0\left[\pi_T^{(\nu)\,\frac{\gamma-1}{\gamma}}\right]^{\gamma}. \tag{9}$$

It is easy to see that for any admissible choice of $\nu \in \mathbf{D}$, the value function in (9) gives an upper bound for the optimal value function of the original problem. In the fictitious market, the wealth dynamics of the portfolio are given by

$$dW_t^{(\nu)} = W_t^{(\nu)}\left[\left(r_t^{(\nu)} + \theta_t^\top \Sigma_{Pt}\eta_t^{(\nu)}\right) dt + \theta_t^\top \Sigma_{Pt}\, dB_t\right] \tag{10}$$

so that

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

The last expression is non-negative according to (4) since $\theta_t \in \mathbf{K}$. Thus, $W_t^{(\nu)} \geq W_t\ \forall\ t \in [0, T]$ and

$$V_0^{(\nu)} \geq V_0. \tag{11}$$

6

In this paper, we use this property of the fictitious problem $(\mathcal{P}^{(\nu)})$ to construct an upper bound on the value function of the original problem $(\mathcal{P})$.

Results in Schroder and Skiadas (2003) imply that if the original optimization problem has a solution, then the upper bound is "tight," i.e., the value function of the fictitious problem $(\mathcal{P}^{(\nu)})$ coincides with the value function of the original problem $(\mathcal{P})$:

$$V_0^{(\nu^*)} \equiv \inf_{\{\nu\}} V^{(\nu)} = V_0. \tag{12}$$

(see Schroder and Skiadas, 2003, Proposition 3(b) and Theorems 7 and 9). The above equality holds for all times, and not just at time 0, i.e., $V_t^{(\nu^*)} = V_t$. Cvitanic and Karatzas (1992) have shown that the solution to the original problem exists under additional restrictions on the utility function, most importantly that the relative risk aversion does not exceed one. Cuoco (1997) proves a more general existence result, imposing minimal restrictions on the utility function.

4   The Performance Bound

The theoretical duality results of Section 3 suggest that one can construct an upper bound on the value function of the portfolio choice problem $(\mathcal{P})$ by computing the value function of any fictitious problem $(\mathcal{P}^{(\nu)})$. The fictitious market is defined by the process ν_t as in (7). Of course, one can pick any fictitious market from the admissible set $\mathbf{D}$ to compute an upper bound. Such a bound is uninformative if it is too loose. Since our objective is to evaluate a particular candidate policy, we can construct a process $\widehat{\nu}_t$ based on such a policy to obtain tighter bounds. The solution to the portfolio choice problem under the fictitious market defined by $\widehat{\nu}_t$ then gives us a performance bound on the candidate policy.

In order to construct the fictitious market as defined by $\widehat{\nu}_t$, we first use the solution to the dual problem (which also gives the solution to the original problem) to establish the link between the optimal policy θ^* and value function $V^{(\nu^*)}$, and the corresponding fictitious asset price processes, as defined by ν^*. Not knowing the optimal portfolio policy and value function, we instead use their approximations to obtain the candidate process for ν^*, which is denoted by $\tilde{\nu}$. This candidate process in general does not belong to $\mathbf{D}$ and cannot be used to define a fictitious problem. Instead, we search for a qualified process $\widehat{\nu}$ in $\mathbf{D}$, which is "closest" to $\tilde{\nu}$. We then use $\widehat{\nu}$ as an approximation to ν^* to define the fictitious problem in $(\mathcal{P}^{(\nu)})$. Since $\widehat{\nu} \in \mathbf{D}$, the solution to the corresponding unconstrained problem in $M^{(\widehat{\nu})}$ provides a valid performance bound for the candidate policy.

According to Cox and Huang (1989),

$$V_t = V_t^{(\nu^*)} = \sup_{\{W_T\}} \; \mathrm{E_t}\left[U(W_T)\right] \quad \text{subject to} \quad \mathrm{E_t}\left[\frac{\pi_T^{(\nu^*)}}{\pi_t^{(\nu^*)}} W_T\right] \leq W_t.$$

The first-order optimality conditions then imply $U'(W_T) = \lambda_t \pi_T^{(\nu^*)}/\pi_t^{(\nu^*)}$ where λ_t is the Lagrange multiplier on the time-t budget constraint. On the other hand, according to the envelope condition, the partial derivative of the value function with respect to the portfolio value satisfies $\partial V_t/\partial W_t = \lambda_t$ (see Karatzas and Shreve 1998, Section 3.7, Theorem 7.7 for a formal proof). Thus, $\pi_T^{(\nu^*)}/\pi_t^{(\nu^*)} =$

$U'(W_T)/(\partial V_t/\partial W_t)$ and hence

$$\frac{\pi_s^{(\nu^*)}}{\pi_t^{(\nu^*)}} = \frac{\partial V_s/\partial W_s}{\partial V_t/\partial W_t}, \quad \forall\, s \geq t.$$

The above equality implies that

$$d\ln\pi_t^{(\nu^*)} = d\ln\frac{\partial V_t}{\partial W_t}. \tag{13}$$

In particular, the stochastic part of $d\ln\pi_t^{(\nu^*)}$ is equal to the stochastic part of $d\ln\partial V_t/\partial W_t$. If V_t is smooth, Ito's lemma, and equations (8) and (2) imply that

$$\eta_t^* \equiv \eta_t^{(\nu^*)} = -W_t\left(\frac{\partial^2 V_t/\partial W_t^2}{\partial V_t/\partial W_t}\right)\Sigma_{Pt}^\top\theta_t^* - \left(\frac{\partial V_t}{\partial W_t}\right)^{-1}\Sigma_{Xt}^\top\left(\frac{\partial^2 V_t}{\partial W_t\partial X_t}\right) \tag{14}$$

where θ_t^* denotes the optimal portfolio policy for the original problem. In the special but important case of a CRRA utility function (which we analyze in more detail below) the expression for $\eta_t^{(\nu^*)}$ simplifies. In particular, the first term in (14) becomes equal to $\gamma\Sigma_{Pt}^\top\theta_t^*$, where γ is the relative risk aversion coefficient of the utility function, and one only needs to compute the first derivative of the value function with respect to the state variables X_t to evaluate the second term in (14). This simplifies numerical implementation, since it is generally easier to estimate first-order than second-order partial derivatives of the value function.

Given an approximation to the optimal portfolio policy $\widetilde{\theta}_t$, one can compute the corresponding approximation to the value function, $\widetilde{V}_t$, defined as the conditional expectation of the utility of terminal wealth, under the portfolio policy $\widetilde{\theta}_t$. We can then construct a process $\widehat{\nu}$ as an approximation to ν^*, using (14). Approximations to the portfolio policy and value function can be obtained using a variety of methods (e.g., Brandt, Goyal, Santa-Clara, 2001). In this paper, we take $\widetilde{\theta}_t$ as given and use it to construct an upper bound on the unknown true value function V_0.

Assuming that the approximate value function $\widetilde{V}$ is sufficiently smooth, we can replace V_t and θ_t^* in (14) with $\widetilde{V}_t$ and $\widetilde{\theta}_t$ and obtain

$$\widetilde{\eta}_t \equiv \eta_t^{(\widetilde{\nu})} = -W_t\left(\frac{\partial^2\widetilde{V}_t/\partial W_t^2}{\partial\widetilde{V}_t/\partial W_t}\right)\Sigma_{Pt}^\top\widetilde{\theta}_t - \left(\frac{\partial\widetilde{V}_t}{\partial W_t}\right)^{-1}\Sigma_{Xt}^\top\left(\frac{\partial^2\widetilde{V}_t}{\partial W_t\partial X_t}\right). \tag{15}$$

We then define $\widetilde{\nu}_t$ as a solution to (7b).

Obviously, $\widetilde{\eta}_t$ is a candidate for the market price of risk in the fictitious market. However, there is no guarantee that $\widetilde{\eta}_t$ and the corresponding process $\widetilde{\nu}_t$ belong to the feasible set $\mathbf{D}$ defined by (6). In fact, as we illustrate below, for many important classes of problems the support function $\delta(\nu_t)$ may be infinite for some values of its argument. Thus, we look for a price-of-risk process $\widehat{\eta}_t \in \mathbf{D}$ that is "close" to $\widetilde{\eta}_t$. We choose a Euclidian norm as our measure of distance between the two processes to make the resulting optimization problem tractable.

To guarantee that $\widehat{\eta}_t \in \mathbf{D}$ without sacrificing tractability, we replace the original integrability

constraints defining the set $\mathbf{D}$ with a set of tighter uniform bounds,

$$\|\widehat{\eta} - \eta\| \leq A_1 \tag{16a}$$

$$\delta(\widehat{\nu}) \leq A_2 \tag{16b}$$

where A_1 and A_2 are positive constants that can be taken to be arbitrarily large. The condition (16a) implies that the process $\widehat{\nu}_t$ is square-integrable, since we have assumed that η_t is square integrable and $\|\widehat{\eta} - \eta\|^2 = \widehat{\nu}^\top \left(\Sigma_P^{-1}\right)^\top \Sigma_P^{-1}\widehat{\nu} \geq A\|\widehat{\nu}\|^2$ for some $A > 0$. If the process η_t^*, corresponding to the optimal portfolio policy, does not satisfy the constraints (16a, 16b), imposing such constraints on $\widehat{\eta}_t$ and $\widehat{\nu}_t$ would increase the upper bound. This is the price one must pay for a computationally tractable description of a feasible set of $\widehat{\eta}_t$ and $\widehat{\nu}_t$. For the examples we consider below, one can verify directly that the constructed processes $\widehat{\eta}_t$ and $\widehat{\nu}_t$ are feasible and therefore constraints (16a, 16b) do not need to be imposed. In general, since the choice of constants A_1 and A_2 is somewhat arbitrary, one may want to experiment with different values. Constraints (16a, 16b) may bind because the optimal process η_T^* is not uniformly bounded, or because the approximation error $\widetilde{\eta}_t - \eta_t^*$ is not uniformly bounded. In the former case, one would like to make the constants as large as possible, while in the latter case large values of A_1 and A_2 could increase the upper bound.

We now define $\widehat{\eta}_t$ and $\widehat{\nu}_t$ as a solution of the following problem:

$$\min_{\widehat{\nu}, \widehat{\eta}} \|\widehat{\eta} - \widetilde{\eta}\|^2 \tag{17}$$

subject to

$$\widehat{\eta} = \eta + \Sigma_P^{-1}\widehat{\nu}, \tag{18a}$$

$$\delta(\nu) < \infty, \tag{18b}$$

$$\|\widehat{\eta} - \eta\| \leq A_1, \tag{18c}$$

$$\delta(\widehat{\nu}) \leq A_2, \tag{18d}$$

The objective is to minimize the Euclidian distance between the process $\widehat{\eta}$, used to compute the upper bound, and the candidate process for the market price of risk, $\widetilde{\eta}_t$. The first constraint simply relates the market price of risk to the process $\widehat{\nu}$. The second constraint requires the support function to be finite. This is a general form of a constraint, which, depending on the specific choice of the set $\mathbf{K}$, represents a set of constraints on $\widehat{\nu}$. For the examples we consider below, the requirement that $\delta(\nu) < \infty$ results in a set of linear constraints on $\widehat{\nu}$. Finally, the last two constraints in (18) are the uniform bounds guaranteeing that $\widehat{\eta} \in \mathbf{D}$, as we discussed above.

The value of $\widehat{\eta}_t$ and $\widehat{\nu}$ can be computed quite easily for many important classes of portfolio choice problems. In particular, we consider the three examples mentioned in Section 2.

Incomplete markets

Assume that only the first L stocks are traded. The positions in the remaining $N - L$ stocks are restricted to zero. In this case the set of feasible portfolio policies is given by

$$\mathbf{K} = \{\theta \mid \theta_i = 0 \text{ for } L < i \leq N\} \tag{19}$$

and hence the support function $\delta(\nu)$ is equal to zero if $\nu_i = 0$, $1 \leq i \leq L$ and is infinite otherwise. Thus, as long as $\nu_i = 0$, $1 \leq i \leq L$, the constraint (16b) does not need to be imposed explicitly. To find $\widehat{\eta}$ and $\widehat{\nu}$ we must solve

$$\min_{\widehat{\eta}, \widehat{\nu}} \|\widehat{\eta} - \widetilde{\eta}\|^2 \tag{20}$$

subject to

$$
\begin{aligned}
\widehat{\eta} &= \eta + \Sigma_P^{-1}\widehat{\nu} \\
\widehat{\nu}_i &= 0, \quad 1 \leq i \leq L \\
\|\widehat{\eta} - \eta\|^2 &\leq A_1^2
\end{aligned}
$$

The diffusion matrix Σ_P is lower triangular and so is its inverse. Using this, the solution can be expressed explicitly as

$$
\begin{aligned}
\widehat{\eta}_i &= \eta_i, \quad 1 \leq i \leq L \\
\widehat{\eta}_j &= \eta_j + a(\widetilde{\eta}_j - \eta_j), \quad L < j \leq N \\
\widehat{\nu} &= \Sigma_P(\widehat{\eta} - \eta)
\end{aligned}
$$

where

$$a = \min\left[1, \left(\frac{A_1^2 - \|\widetilde{\eta} - \eta\|_{(L)}^2}{\|\widetilde{\eta} - \eta\|^2 - \|\widetilde{\eta} - \eta\|_{(L)}^2}\right)^{1/2}\right], \quad \|\widetilde{\eta}\|_{(L)}^2 = \sum_{i=1}^{L} \widetilde{\eta}_i^2.$$

Note that when the constant A_1 is very large, it will invariably be the case that $a = 1$, and therefore $\widehat{\eta}_j = \widetilde{\eta}_j$ for $L < j \leq N$.

Incomplete markets and no short sales

Consider the market in which only the first L stocks can be traded and no short sales (of the stocks) are allowed. The set of admissible portfolios is given by

$$\mathbf{K} = \{\theta \mid \theta \geq 0, \ \theta_i = 0 \text{ for } L < i \leq N\} \tag{21}$$

and the support function is equal to zero if $\nu_i \geq 0$, $i = 1, \ldots, L$ and is infinite otherwise. As in the previous case, the constraint (16b) is automatically satisfied. Thus, $\widehat{\nu}$ can be determined as a solution of a standard quadratic programming problem:

$$\min_{\widehat{\nu}, \widehat{\eta}} \|\widehat{\eta} - \widetilde{\eta}\|^2 \tag{22}$$

10

subject to

$$\begin{aligned}
\widehat{\eta} &= \eta + \Sigma_P^{-1}\widehat{\nu} \\
\widehat{\nu}_i &\geq 0, \quad 1 \leq i \leq L \\
\|\widehat{\eta} - \eta\|^2 &\leq A_1^2.
\end{aligned}$$

Incomplete markets, no short sales, and no borrowing

Consider the same market as in the previous case, but no short sales and borrowing are allowed. Then the set of admissible portfolios is given by

$$\mathbf{K} = \{\theta \mid \theta \geq 0,\ 1^\top \theta \leq 1,\ \theta_i = 0 \text{ for } L < i \leq N\}. \tag{23}$$

The support function is given by $\delta(\nu) = \max(0, -\nu_1, \ldots, -\nu_L)$, which is finite for any vector ν. Because in this case $\delta(\nu) \leq \|\nu\|$, the relation $\|\nu\| = \|\Sigma_P(\widehat{\eta} - \eta)\| \leq \|\Sigma_P\|A_1$ (see 16a, 16b) implies that as long as A_2 is sufficiently large compared to A_1, one only needs to impose (16a) and (16b) is redundant. Thus, one needs to solve the following problem:

$$\min_{\widehat{\eta}, \widehat{\nu}} \ \|\widehat{\eta} - \widetilde{\eta}\|^2 \tag{24}$$

subject to

$$\begin{aligned}
\widehat{\eta} &= \eta + \Sigma_P^{-1}\widehat{\nu} \\
\|\widehat{\eta} - \eta\|^2 &\leq A_1^2.
\end{aligned}$$

Then the fictitious market is described by

$$\begin{aligned}
\widehat{\eta} &= \eta + \min\left(1, \frac{A_1}{\|\widetilde{\eta} - \eta\|}\right)(\widetilde{\eta} - \eta) \\
\widehat{\nu} &= \Sigma_P(\widehat{\eta} - \eta).
\end{aligned}$$

Again, by setting A_1 to be sufficiently large we will invariably obtain $\widehat{\eta} = \widetilde{\eta}$.

5 Numerical Experiments

The method described in the previous section can be applied to various problems of portfolio choice. In this section we illustrate its performance based on several numerical experiments. But first we summarize our algorithm as a sequence of four basic steps:

1. Start with an approximation to the optimal portfolio policy of the original problem, and the corresponding approximation to the value function, which can be obtained using a variety of methods. In this paper we simply assume that such approximations are available.

2. Use the approximate portfolio policy and partial derivatives of the approximate value function to construct a process $\widetilde{\eta}_t$ according to the explicit formula (15). The process $\widetilde{\eta}_t$ is a candidate for the market price of risk in the fictitious market.

3. Construct a process $\widehat{\eta}_t$ that is close to $\widetilde{\eta}_t$ and satisfies the conditions for the market price risk of a fictitious market in the dual problem. This involves solving a quadratic optimization problem (17).

4. Compute the value function from the static problem $(\mathcal{P}^{(\nu)})$ in the resulting fictitious market defined by the market price of risk process $\widehat{\eta}_t$. This can be accomplished efficiently using Monte Carlo simulation. This results in an upper bound on the value function of the original problem.

The lower bound on the value function is obtained by simulating the terminal wealth distribution under the approximate portfolio strategy.

Calibration

In our numerical experiments we assume that the utility function is of the constant relative risk aversion (CRRA) type so that $U(W) = W^{1-\gamma}/(1-\gamma)$. We consider three values for the relative risk aversion parameter $\gamma = 1.5$, 3, and 5. We also consider two values for the problem horizon: $T = 5$ and $T = 10$ years.

Our continuous-time model is very similar to the discrete-time market model in Lynch (2001) and we use the parameters estimated by Lynch to guide our calibration. He considers a financial market with three traded risky assets and a single state variable. To represent the same market in our framework, we assume that there is a total of four risky assets, the first three of which are traded, and a single state variable, i.e., we consider an incomplete market with $N = 4$, $M = 1$, and $L = 3$, meaning that the first three stocks can be traded.

The dynamics of asset returns is given by

$$r_t = r \tag{25a}$$

$$dP_t = P_t\left[(\mu_0 + X_t\mu_1)\,dt + \Sigma_P\,dB_t\right] \tag{25b}$$

$$dX_t = -kX_t\,dt + \Sigma_X\,dB_t \tag{25c}$$

The diffusion matrix Σ_X is of size 1 by 4 and we assume that the fourth row of the matrix Σ_P coincides with Σ_X. This assumption is made for convenience, and it economizes on the number of parameters to be reported. Note that since the fourth risky asset is not traded, one can specify the fourth row of Σ_P arbitrarily, as long as the matrix Σ_P satisfies the necessary regularity conditions. We set the initial value of the state variable to zero, $X_0 = 0$, in all numerical examples below.

Lynch considers two choices for the state variable X: the dividend yield and the term spread. The dividend yield captures the rate at which dividends are paid out, as a fraction of the total stock market value. Specifically, Lynch uses a continuously compounded twelve-month yield on the value-weighted NYSE index. The term spread is the difference in yields between twenty-year and one-month Treasury securities. Both of these predictive variables are normalized to have zero mean and unit variance. He also considers two sets of risky assets: three portfolios obtained by sorting stocks on their size or on their book-to-market ratio. Details of the procedure are reported in Lynch

(2001), here we are interested only in the parameters of the stock return process, as summarized in Tables 1 and 2 of Lynch (2001). Thus, by considering two choices of risky assets and two choices of the predictive variable, we have four sets of calibrated parameter values. These are reported in Table 1 below. We set the riskfree rate $r = 0.01$ throughout.

Approximate policies

As in Lynch (2001), we consider two types of portfolio policies. The first policy, which we call "static", ignores predictability of stock returns. It is defined using the unconditional average returns instead of the time-varying conditional expected returns on the stocks. Specifically,

$$\widetilde{\theta}_t^{static} = \arg\max_{\theta} \ (\mu_0^\top - r)\theta - \frac{1}{2}\gamma\,\theta^\top \Sigma_P \Sigma_P^\top \theta, \quad \text{subject to } \theta \in \mathbf{K}. \tag{26}$$

This is an optimal portfolio policy in a dynamic model with a constant investment opportunity set and cone constraints on portfolio positions, e.g., (Karatzas and Shreve, 1998, Sec. 6.6).

As we discussed in the introduction, a small gap between the upper bound and the expected utility under the static policy indicates that return predictability is not economically important for the problem at hand. The reverse, however, is not necessarily true — when the upper bound is significantly higher than the lower bound, we cannot formally conclude that the static portfolio policy is significantly inferior to the optimal strategy. We discuss these issues in more detail in the context of the numerical examples below.

The second portfolio policy we consider is the "myopic" policy. It is defined in the same way as a static policy, except the instantaneous moments of asset returns are fixed at their current values, as opposed to their long-run average values:

$$\widetilde{\theta}_t^{myopic} = \arg\max_{\theta} \ (\mu_{Pt}^\top - r)\theta - \frac{1}{2}\gamma\,\theta^\top \Sigma_P \Sigma_P^\top \theta, \quad \text{subject to } \theta \in \mathbf{K}. \tag{27}$$

The approximate portfolio policy in (27) ignores the hedging component of the optimal trading strategy. By comparing the corresponding expected utility with an upper bound we will be able to evaluate the economic importance of hedging as a component of the portfolio policy. In a portfolio choice problem similar to ours, Lynch finds a significant hedging demand for some of the parameter values he considers. In particular, for the third parameter set when there are no portfolio constraints, Lynch finds that the optimal policy is very different from the myopic approximation. Below we will show that even though the myopic strategy may be far from optimal in terms of actual portfolio holdings, the utility loss due to using it may not be very large. As we previously emphasized, the advantage of our method is that such results can be obtained without knowing the optimal strategy, which would be the case for most realistic applications.

Our algorithm for evaluating portfolio strategies relies on an approximate value function $\widetilde{V}_t$. In general, one has to estimate $\widetilde{V}_t$ as the conditional expectation of the utility of terminal wealth under the approximate strategy $\widetilde{\theta}_t$. This could be done for example using standard regression methods. Instead, we choose a particularly straightforward closed-form expression for $\widetilde{V}_t$. This simplifies computations, while still allowing us to illustrate the conceptual steps involved in estimating the upper bound on the value function. The cost of using a simplified closed-form formula for $\widetilde{V}_t$ is a potentially wider gap between the upper bound and the true value function.

Because the utility function is of CRRA type, the true value function has a homothetic functional form, $V_t = g(t, X_t) W_t^{1-\gamma}$. As a drastic simplification, we ignore the dependence of the value function on the state variable and set $\widetilde{V}_t = \widetilde{g}(t) W_t^{1-\gamma}$. The exact form of the function $\widetilde{g}(t)$ does not affect the specification of the artificial market below. Note that removing the dependence of the approximate value function on X does not eliminate it from the price of risk in the fictitious market, given by (14). This is because the leading term in that expression depends on X_t through the approximate portfolio policy, $\widetilde{\theta}_t$.

We will see in the numerical examples below, even such a simplistic approximation to the value function can lead to informative bounds on the true value function. However, in realistic practical applications, one should use more accurate approximations to the value function to achieve tighter bounds whenever possible.

Simulation

For each of the approximate portfolio strategies, we estimate the corresponding expected utility of terminal wealth as well as the upper bound on the true value function using Monte Carlo simulation. The expected utility under the suboptimal portfolio policy provides a lower bound on the true value function of the problem. We estimate the former by simulating independently $100,000$ or $1,000,000$ paths of stock prices and state variables according to (25a)–(25c). We used $100,000$ paths to estimate the bounds for the problems with constraints on short sales only as the corresponding quadratic subproblem in that case needs to be solved numerically. We used $1,000,000$ paths for all other cases. (The problems were solved using *Matlab* so the computational times for the various problems are not particularly informative.) We compute the portfolio value along each path under the approximate portfolio strategy and record the utility of the terminal wealth. We then average over the simulated paths to estimate the expected utility. In simulations, we discretize the continuous-time diffusion processes using a standard Euler scheme with a time step $\triangle t = 0.01$ years. For the case of incomplete markets and the case of no short sales and no borrowing, quadratic subproblems given by eqns (17) and (18) can be solved in closed form. For the case of no borrowing, we solved each of the quadratic subproblems using the MATLAB optimizer. Each of the subproblems was in three dimensions with three inequality constraints. The total number of subproblems was equal to $10^7 \times T$.

As we showed in Section 4, the true value function is bounded from above by the value function in the fictitious market, which in turn is given by

$$V_0^{(\nu)} \;=\; \frac{W_0^{1-\gamma}}{1 - \gamma} \; \mathrm{E}_0 \left[\pi_T^{(\nu) \frac{\gamma-1}{\gamma}} \right]^{\gamma} \tag{28}$$

We simulate the process $\pi_T^{(\nu)}$ by discretizing the integrals in (8) with the same time step $\triangle t = 0.01$. The market price of risk and the riskfree rate in the fictitious market are computed according to the procedure described in Section 4. In our simulations we do not impose the constraint (16a), since one can directly verify under our assumption about the return generating process and the approximate trading strategy that the process $\widehat{\nu}_t$ is square integrable. The second constraint (16b) does not need to be imposed explicitly for the problems we are considering below, as discussed in Section 4.

Incomplete Markets

We first consider the incomplete market model, in which portfolio positions in tradable asset are unconstrained. Table 2 reports the estimates of the expected utility both under the static and myopic portfolio strategies in equations (26) and (27), which provide lower bounds on the true value function, as well as the corresponding estimates of the upper bound. Expected utility is reported as a continuously compounded certainty equivalent return, R. The value of R corresponding to the value function V_0 is defined by $U\left(W_0 e^{RT}\right) = V_0$. Alternatively, one could report the utility-equivalent wealth W, defined by $U\left(W\right) = V_0$. We express our results in terms of the certainty equivalent return to facilitate comparison across different problem horizons T. We also report the exact value function, which can be computed explicitly using the results in Kim and Omberg (1996) and Liu (2005). Table 2 contains results for four parameters sets, three values of the risk aversion parameter γ and two values of the economy horizon T.

One can see that there is a sizeable gap between the expected utility achieved by the static strategy and the corresponding upper bound on the true value function. For most of the parameter combinations considered, the difference, expressed as annualized certainty equivalent return, is in the range of $1\% - 3\%$.

The difference between the upper bound and the lower bound can be decomposed into a sum of two terms:

$$UB - LB = (V - LB) + (UB - V). \tag{29}$$

The first term is the utility loss of the approximate portfolio strategy under consideration. It captures the utility loss due to using a sub-optimal portfolio strategy. The second term is the difference between the upper bound and the true value function and it depends on how tight the estimated upper bound is. Ideally, we would like to know the first term, but since the true value function is not known, we can only estimate the sum of two terms, i.e., the difference $UB - LB$. This difference is large because either or both of the two terms on the right-hand side are large. In particular, the magnitude of the second term depends on the accuracy of the approximate portfolio policy and value function used to estimate the upper bound. Table 2 indicates that in this particular case, the magnitude of the second term is relatively small, as can be seen from the difference between the upper bound, UB^s, and the exact solution, V^u. (The exact solution is available for this special case from Liu (2005). This will not be the case once we impose portfolio constraints below.) This is somewhat surprising, since the difference between the lower and the upper bound suggests that the static policy is quite far from being optimal.

To summarize, the above results do not allow us to conclude that predictability can be safely ignored for the problems under consideration. However, without knowing the exact solutions we could not argue that the economic value of predictability is large either. More generally, one should not view the difference $UB - LB$ we construct as a formal test statistic for optimality of a given portfolio policy, since we haven't established that the difference $UB - LB$ must be "small" when a portfolio policy used to construct the upper bound is close to being optimal. One can, however, draw definitive conclusions from the cases when the difference $UB - LB$ is small, which implies that the considered suboptimal portfolio policy has a small utility cost.

We now turn to the corresponding results for the myopic strategy, also shown in Table 2. As we have discussed in the introduction, using our method, one can verify that for a particular problem

under consideration ignoring hedging component of demand has little effect on the expected utility. For the problems with a five-year horizon, the difference between the point estimates of UB and LB tends to be in the .1% – .2% range, except for the third parameter set, where it reaches approximately .6%. Thus, we conclude that for most of the parameter sets, the hedging component of demand is not very significant economically with a five-year model horizon.

When we increase the problem horizon to ten years, the gap $UB - LB$ increases. For the second and fourth parameter sets the difference remains small, under .2%, but for the third parameter set it now reaches as much as 1.5%, for the case of $\gamma = 5$. These observations are consistent with our intuition that hedging demand tends to be more important for problems with longer horizons. It is well known that the myopic policy is optimal for $\gamma = 1$ in this case, and therefore it is not surprising that it performs well for values of γ close to one.

No short sales and no borrowing

We now consider the market in which both short sales and borrowing are prohibited, i.e., portfolio policies must satisfy (23). Table 3 reports the estimates of the expected utility both under the static and myopic portfolio strategies and the corresponding estimates of the upper bound. We also report the value function obtained by relaxing the portfolio constraints on the three risky assets, which is the same as the values reported in Table 2. Table 3 contains results for four parameters sets, three values of the risk aversion parameter γ and two values of the economy horizon T.

As in the case of incomplete markets and no position constraints, there is a significant gap between the expected utility achieved by the static strategy and the corresponding upper bound on the true value function. As before, we decompose the difference between the upper bound and the lower bound as in (29). Table 3 indicates that the magnitude of the second term is substantial at least for some of the parameter sets. In particular, for $\gamma = 5$, the estimated upper bound often exceeds the value function of the unconstrained problem, which in turn is greater then or equal to the true value function. As in the incomplete market case, we cannot conclude that predictability can be safely ignored for the problems under consideration.

The results for the myopic strategy are similar to those in Table 2. For the problems with a five-year horizon, the difference between the point estimates of UB and LB tends to be in the .1% – .2% range, except for the third parameter set, where it reaches approximately .5%. Thus, we conclude that for most of the parameter sets, the hedging component of demand is not very significant economically with a five-year model horizon. A closer examination of the portfolio composition under the myopic policy (not reported here) suggests that our conclusions may be too conservative, at least for the third parameter set. Because of the binding short sale constraints, the myopic policy is likely to be close to the optimum most of the time. Thus, if we were to use a more accurate approximation to the approximate value function $\widetilde{V}_t$, our upper bound would likely be tighter than the values reported in Table 3. As we pointed out earlier, our objective here is not to maximize the accuracy of computations, but rather to illustrate the concepts, and we therefore do not pursue this issue any further.

No short sales

In the numerical experiments above we assumed that neither borrowing nor short sales are allowed. As we showed in Section 4, in this case the fictitious market used to estimate the upper bound is particularly straightforward to define. We now consider the case when short sales are prohibited, but borrowing is allowed. Then, to define the fictitious market one must solve a quadratic optimization problem (22).

Table 4 reports the results and in order to reduce the computational effort we only consider the case of $T = 5$ years. Unlike the previous two cases, when we impose no short sales constraint we have to solve a quadratic optimization problem numerically for every point in time on each simulated path. This increases the computational cost. To reduce the total computation time, we limit the model horizon to $T = 5$ years and only use $100,000$ simulated paths.) A comparison with Table 3 reveals a peculiar regularity. Because the borrowing constraint rarely binds for the myopic policy when $\gamma = 3$ and $\gamma = 5$, the corresponding expected utility, LB, is similar to the values reported in Table 3. However, the upper bound tends to be much tighter for these cases compared to Table 3, even though the true value function for the problem considered in Table 4 exceeds the one for the problem corresponding to Table 3 (since the latter problem is a more constrained version of the former). The situation is reversed for $\gamma = 1.5$, but in that case the static policy calls for borrowing, hence the expected utilities cannot be directly compared across the two tables. Nevertheless, one can observe that for $\gamma = 1.5$ the difference $UB - LB$ is significantly larger than for $\gamma = 3$ and $\gamma = 5$.

For the myopic strategy, the difference $UB - LB$ is very similar to the corresponding values in Table 3. It is as high as .4% for the parameter set three and is less than .2% for all other parameter sets, demonstrating that the hedging component of demand is not crucial when borrowing is allowed. A comparison with the results for the static strategy shows that the gap $UB - LB$ for the static strategy is primarily due to the difference between the true value function and the lower bound, and the corresponding upper bounds are relatively tight.

6 Conclusion

In this paper we developed a method for evaluating the quality of approximations for portfolio choice problems by providing an upper bound on the utility loss associated with an approximate solution. Our algorithm relies on the duality theory for constrained portfolio optimization. Starting from an approximate portfolio policy, we construct a fictitious financial market without portfolio constraints in which the optimization problem can be easily solved. The resulting value function is estimated using Monte Carlo simulation and provides an upper bound on the (unknown) value function of the original problem.

Our method is independent of the nature of the original approximation and can be used in combination with various algorithms for approximating optimal portfolio strategies and value functions. The algorithm can also be applied to verify that an approximate strategy is in fact close to the optimum in terms of expected utility. This applies not only to problems that can be potentially handled by traditional numerical methods, but also to large-scale problems, since the computational time required by our simulation-based method does not grow exponentially with the dimension of the state space or the number of risky assets.

17

In our analysis we consider problems with the objective defined over the terminal wealth. Extending our results to models with intermediate consumption is straightforward and will require an approximate consumption policy in addition to the trading strategy. Certain consumption constraints can be handled as well, for instance constraints of the form $C_t \geq \underline{C}$. Such constraints allow the martingale methods to be applied to the dynamic portfolio-consumption choice (e.g., Cox and Huang, 1989).

Our formulation currently does not cover an important class of portfolio choice problems with illiquid (nontradable) assets and liquidity constraints. Such problems are of great practical and theoretical importance, but they rarely allow for explicit analytical solutions and are mostly handled using numerical methods. Since general results for the quality of such numerical approximations are not available, an algorithm for estimating the quality of approximate portfolio policies would be very valuable.

Another promising direction for future research is to explore combining our methodology with recently developed approximate dynamic programming (ADP) techniques. The bounds we obtain could be used for both guiding the choice of approximation architecture in ADP applications, as well as evaluating the resultant policies.

We are pursuing the above extensions in ongoing research.

Acknowledgements

Ait-Sahalia, Y. and M. Brandt. 2001. Variable Selection for Portfolio Choice. *Journal of Finance.* **56** 1297-1351.

Ang, A. and G. Bekaert. 2002. International Asset Allocation with Time-Varying Correlations. Working paper. Columbia University, New York.

Balduzzi, P. and A. Lynch. 1999. Transaction Costs and Predictability: Some Utility Cost Calculations. *Journal of Financial Economics.* **52** 47-78.

Barberis, N. 2002. Investing for the Long Run When Returns are Predictable. *Journal of Finance.* **55** 225-264.

Brandt, M. 1999. Estimating Portfolio and Consumption Choice: A Conditional Euler Equations Approach. *Journal of Finance.* **54** 1609-1645.

Brandt, M., A. Goyal, and P. Santa-Clara. 2001. A Simulation Approach to Dynamic Portfolio Choice with an Application to Industry Momentum. Working paper. UCLA, Los Angeles.

Brennan, M.. 1998. The Role of Learning in Dynamic Portfolio Decisions. *European Finance Review.* **1** 295-306.

Brennan, M. and Y. Xia. 2001. Dynamic Asset Allocation Under Inflation. Working paper. UCLA, Los Angeles.

Brennan, M., E. Schwartz, and R. Lagnado. 1997. Strategic Asset Allocation. *Journal of Economic Dynamics and Control.* **21** 1377-1403.

Campbell, J. and L. Viceira. 1999. Consumption and Portfolio Decisions When Expected Returns are Time Varying. *The Quarterly Journal of Economics.* **114** 433-495.

Chan, Y. and L. Kogan. 2001. Catching Up with the Joneses: Heterogeneous Preferences and the Dynamics of Asset Prices. Forthcoming in *Journal of Political Economy.*

Campbell, J. and L. Viceira. 2002. *Strategic Asset Allocation – Portfolio Choice for Long-Term Investors.* Oxford University Press, USA.

Campbell, J., Y. Chan and L. Viceira. 2002. A Multivariate Model of Strategic Asset Allocation. Forthcoming in *Journal of Financial Economics.*

Chacko, G. and L. Viceira. 2002. Dynamic Consumption and Portfolio Choice with Stochastic Volatility in Incomplete Markets. Working paper, Harvard University, Boston.

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.

Cvitanic, J. and I. Karatzas. 1992. Convex Duality in Constrained Portfolio Optimization. *Annals of Applied Probability.* **2** 767-818.

Dammon, R., C. Spatt and H. Zhang. 2000. Optimal Consumption and Investment with Capital Gains Taxes. *Review of Financial Studies.* **14** 583-616.

Detemple, J., R. Garcia, and M. Rindisbacher. 2000. A Monte-Carlo Method for Optimal Portfolios, CIRANO Working Paper 2000s-05.

Gomes, F. 2002. Exploiting Short-Run Predictability. Working paper. London Business School, London.

Haugh, M. and L. Kogan. 2003. Pricing American Options: A Duality Approach. Forthcoming in *Operations Research*.

den Haan, W. and A. Marcet. 1994. Accuracy in Simulations. *Review of Economic Studies.* **61** 3-17.

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.* **52** 259-304.

Johannes, M., N. Polson and J. Stroud. 2002. Sequential Optimal Portfolio Performance: Market and Volatility Timing. Working paper. Columbia University, New York.

Judd, K. 1996. Approximation, Perturbation, and Projection Methods in Economic Analysis. In H.M. Amman, D.A. Kendrick, J. Rust (eds.), *Handbook of Computational Economics, Volume I.* Elsevier Schience, Amsterdam.

Kandel, S. and R. Stambaugh. On the Predictability of Stock Returns: An Asset-Allocation Perspective. *Journal of Finance.* **51** 385-424.

Karatzas, I., J. Lehoczky, S. Sethi and S. Shreve. 1986. Explicit Solution of a General Consumption/Investment Problem. *Mathematics of Operationas Research.* **11** 261-294.

Karatzas, I., J. Lehoczky, S. Shreve, and G. Xu. 1991. Martingale and Duality Methods for Utility Maximization in an Incomplete Market. *SIAM Journal of Control and Optimization.* **29** 702-730.

Karatzas, I. and S. Shreve. 1998. *Methods of Mathematical Finance.* Springer-Verlag, New York, NY.

Kim, T. S. and E. Omberg. 1986. Dynamic Nonmyopic Portfolio Behavior. *Review of Financial Studies.* **9** 141-61.

Munk, C. 1998. Optimal Consumption/Investment Policies with Undiversifiable Income Risk and Liquidity Constraints. Working paper. Odense University, Denmark.

Liu, J. 2005. Portfolio Choice In Stochastic Environments. Working paper. Stanford University, Palo Alto.

Lynch, A. 2001. Portfolio Choice and Equity Characteristics: Characterizing the Hedging Demands Induced by Return Predictability. *Journal of Financial Economics.* **62** 67-130.

Pastor, L. 2000. Portfolio Selection and Asset Pricing Models. *Journal of Finance.* **55** 179-223.

Pastor, L. and R. Stambaugh. 2000. Comparing Asset Pricing Models: An Investment Perspective. *Journal of Financial Economics.* **56** 335-81.

Pliska, S. 1986. A stochastic Calculus Model of Continuous Trading: Optimal Portfolios. *Mathematics of Operations Research.* **11** 371-382.

Pliska, S. 1997. Introduction to Mathematical Finance: Discrete Time Models. S. Pliska. Blackwell Publishers Inc., Malden, MA.

Rust, J., 1996. Approximation, Perturbation, and Projection Methods in Economic Analysis. In H.M. Amman, D.A. Kendrick, J. Rust (eds.), *Handbook of Computational Economics, Volume I.* Elsevier Schience, Amsterdam.

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, IL.

Stambaugh, R. 1999. Predictive Regressions. *Journal of Financial Economics.* **54** 375-421.

Xia, Y. 2001. Learning About Predictability: The Effect of Parameter Uncertainty on Optimal Dynamic Asset Allocation. *Journal of Finance* **56**. 205-247.

Parameter set 1

k	μ_0	μ_1	Σ_P			
0.366	0.081	0.034	0.186	0.000	0.000	0.000
	0.110	0.059	0.228	0.083	0.000	0.000
	0.130	0.073	0.251	0.139	0.069	0.000
	0.000	0.000	-0.741	-0.037	-0.060	0.284

Parameter set 2

k	μ_0	μ_1	Σ_P			
1.671	0.081	0.046	0.186	0.000	0.000	0.000
	0.110	0.070	0.227	0.082	0.000	0.000
	0.130	0.086	0.251	0.139	0.069	0.000
	0.000	0.000	-0.017	0.149	0.058	1.725

Parameter set 3

k	μ_0	μ_1	Σ_P			
0.366	0.142	0.065	0.256	0.000	0.000	0.000
	0.109	0.049	0.217	0.054	0.000	0.000
	0.089	0.049	0.207	0.062	0.062	0.000
	0.000	0.000	-0.741	0.040	0.034	0.288

Parameter set 4

k	μ_0	μ_1	Σ_P			
1.671	0.142	0.061	0.256	0.000	0.000	0.000
	0.109	0.060	0.217	0.054	0.000	0.000
	0.089	0.067	0.206	0.062	0.062	0.000
	0.000	0.000	-0.017	0.212	0.096	1.716

Table 1: **Calibrated model parameters.** The four sets of parameters correspond to: (1) size sorted portfolios and the dividend yield as a state variable; (2) size sorted portfolios and the term spread as a state variable; (3) book-to-market sorted portfolios and the dividend yield as a state variable; (4) book-to-market sorted portfolios and the term spread as a state variable. Parameter values are based on the estimates in Tables 1 and 2 of Lynch (2001).

	$T=5$			$T=10$		
	$\gamma=1.5$	$\gamma=3$	$\gamma=5$	$\gamma=1.5$	$\gamma=3$	$\gamma=5$
			Parameter set 1			
LB^s	7.50	4.64	3.27	7.70	4.84	3.42
	(7.48, 7.52)	(4.63, 4.65)	(3.27, 3.28)	(7.69, 7.71)	(4.84, 4.85)	(3.42, 3.43)
UB^s	9.44	5.94	4.20	10.12	6.72	4.88
	(9.40, 9.48)	(5.86, 6.01)	(4.10, 4.30)	(10.09, 10.15)	(6.65, 6.79)	(4.76, 5.00)
LB^m	9.35	5.72	3.97	9.94	6.18	4.29
	(9.32, 9.37)	(5.70, 5.73)	(3.96, 3.98)	(9.92, 9.96)	(6.16, 6.19)	(4.28, 4.29)
UB^m	9.44	5.94	4.20	10.12	6.72	4.88
	(9.40, 9.48)	(5.86, 6.01)	(4.10, 4.30)	(10.09, 10.15)	(6.65, 6.79)	(4.76, 5.00)
V^u	9.44	5.95	4.19	10.09	6.62	4.75
			Parameter set 2			
LB^s	6.55	3.69	2.59	6.53	3.67	2.58
	(6.52, 6.58)	(3.67, 3.70)	(2.58, 2.60)	(6.51, 6.55)	(3.66, 3.69)	(2.58, 2.59)
UB^s	9.07	5.00	3.40	9.16	5.07	3.46
	(9.02, 9.11)	(4.95, 5.05)	(3.35, 3.46)	(9.13, 9.20)	(5.03, 5.11)	(3.41, 3.51)
LB^m	9.02	4.92	3.33	9.10	4.95	3.35
	(8.99, 9.05)	(4.90, 4.93)	(3.32, 3.34)	(9.08, 9.12)	(4.94, 4.97)	(3.34, 3.36)
UB^m	9.07	5.00	3.40	9.16	5.07	3.46
	(9.02, 9.11)	(4.95, 5.05)	(3.35, 3.46)	(9.13, 9.20)	(5.03, 5.11)	(3.41, 3.51)
V^u	9.02	4.92	3.33	9.08	4.94	3.34

Parameter set 3

LB^s	14.67	8.59	5.73	15.06	8.98	6.02
	(14.64, 14.70)	(8.57, 8.60)	(5.72, 5.74)	(15.04, 15.08)	(8.97, 8.99)	(6.01, 6.02)
UB^s	16.84	10.43	7.27	17.84	11.90	8.76
	(16.78, 16.90)	(10.32, 10.55)	(7.09, 7.45)	(17.80, 17.88)	(11.78, 12.01)	(8.54, 8.98)
LB^m	16.64	9.88	6.61	17.47	10.60	7.10
	(16.61, 16.68)	(9.86, 9.90)	(6.59, 6.62)	(17.44, 17.49)	(10.58, 10.61)	(7.09, 7.12)
UB^m	16.84	10.43	7.27	17.84	11.90	8.76
	(16.78, 16.90)	(10.32, 10.55)	(7.09, 7.45)	(17.80, 17.88)	(11.78, 12.01)	(8.54, 8.98)
V^u	16.79	10.32	7.06	17.76	11.55	8.12

Parameter set 4

LB^s	13.17	7.10	4.66	13.18	7.11	4.67
	(13.14, 13.21)	(7.08, 7.12)	(4.65, 4.68)	(13.15, 13.21)	(7.09, 7.13)	(4.66, 4.69)
UB^s	15.93	8.47	5.54	16.06	8.59	5.62
	(15.87, 15.99)	(8.40, 8.55)	(5.46, 5.63)	(16.01, 16.10)	(8.52, 8.66)	(5.53, 5.71)
LB^m	15.91	8.45	5.48	16.02	8.51	5.51
	(15.87, 15.95)	(8.42, 8.47)	(5.46, 5.49)	(15.99, 16.05)	(8.49, 8.54)	(5.49, 5.53)
UB^m	15.93	8.47	5.54	16.06	8.59	5.62
	(15.87, 15.99)	(8.40, 8.55)	(5.46, 5.63)	(16.01, 16.10)	(8.52, 8.66)	(5.53, 5.71)
V^u	15.91	8.45	5.47	16.00	8.49	5.49

Table 2: **Incomplete markets.** The four parameter sets are defined in Table 1. All results are computed for the initial value of the state variable $X_0 = 0$. The rows marked LB^s and LB^m report the estimates of the expected utility achieved by using the static and myopic portfolio strategies, respectively. The estimates are based on $1,000,000$ independent simulations. Expected utility is reported as a continuously compounded certainty equivalent return. Approximate 95% confidence intervals are reported in parentheses. The rows marked UB^s and UB^m report the analogous results for the upper bound on the true value function computed according to the procedure described in Section 4. The row marked V^u reports the optimal value function for the problem.

	$T = 5$			$T = 10$		
	$\gamma = 1.5$	$\gamma = 3$	$\gamma = 5$	$\gamma = 1.5$	$\gamma = 3$	$\gamma = 5$

<u>Parameter set 1</u>

	$\gamma = 1.5$	$\gamma = 3$	$\gamma = 5$	$\gamma = 1.5$	$\gamma = 3$	$\gamma = 5$
LB^s	7.39	4.64	3.28	7.57	4.84	3.42
	(7.37, 7.41)	(4.63, 4.65)	(3.27, 3.28)	(7.56, 7.58)	(4.83, 4.85)	(3.42, 3.43)
UB^s	9.22	7.23	6.33	9.59	7.80	7.00
	(9.19, 9.25)	(7.19, 7.26)	(6.29, 6.37)	(9.57, 9.62)	(7.78, 7.83)	(6.97, 7.02)
LB^m	8.09	5.57	3.91	8.35	5.99	4.22
	(8.07, 8.11)	(5.56, 5.59)	(3.89, 3.92)	(8.34, 8.36)	(5.98, 6.00)	(4.21, 4.23)
UB^m	8.15	5.84	4.18	8.45	6.48	4.77
	(8.12, 8.18)	(5.79, 5.90)	(4.08, 4.28)	(8.43, 8.47)	(6.42, 6.54)	(4.60, 4.94)
V^u	9.44	5.95	4.19	10.09	6.62	4.75

<u>Parameter set 2</u>

	$\gamma = 1.5$	$\gamma = 3$	$\gamma = 5$	$\gamma = 1.5$	$\gamma = 3$	$\gamma = 5$
LB^s	6.54	3.69	2.60	6.52	3.67	2.58
	(6.52, 6.57)	(3.68, 3.71)	(2.59, 2.61)	(6.51, 6.54)	(3.66, 3.69)	(2.58, 2.59)
UB^s	9.54	6.94	5.79	9.59	7.00	5.84
	(9.50, 9.57)	(6.90, 6.98)	(5.74, 5.83)	(9.57, 9.62)	(6.96, 7.03)	(5.80, 5.87)
LB^m	7.88	4.85	3.30	7.91	4.87	3.31
	(7.86, 7.90)	(4.84, 4.87)	(3.29, 3.31)	(7.90, 7.93)	(4.86, 4.88)	(3.30, 3.32)
UB^m	7.93	4.98	3.42	7.95	5.01	3.44
	(7.89, 7.96)	(4.93, 5.03)	(3.37, 3.48)	(7.93, 7.98)	(4.97, 5.04)	(3.39, 3.49)
V^u	9.02	4.92	3.33	9.08	4.94	3.34

Parameter set 3

LB^s	10.03 (10.01, 10.04)	6.76 (6.75, 6.78)	4.62 (4.61, 4.63)	10.22 (10.21, 10.23)	7.12 (7.11, 7.12)	4.87 (4.87, 4.88)
UB^s	10.71 (10.68, 10.74)	8.83 (8.78, 8.88)	7.38 (7.33, 7.42)	11.04 (11.02, 11.06)	9.51 (9.47, 9.54)	8.17 (8.13, 8.21)
LB^m	10.17 (10.15, 10.18)	7.65 (7.63, 7.66)	5.35 (5.33, 5.36)	10.37 (10.36, 10.38)	8.18 (8.17, 8.19)	5.80 (5.79, 5.81)
UB^m	10.23 (10.20, 10.26)	8.05 (7.98, 8.11)	5.85 (5.69, 6.01)	10.47 (10.46, 10.49)	8.88 (8.82, 8.94)	6.79 (6.34, 7.25)
V^u	16.79	10.32	7.06	17.76	11.55	8.12

Parameter set 4

LB^s	9.22 (9.20, 9.24)	5.31 (5.29, 5.32)	3.57 (3.56, 3.58)	9.21 (9.19, 9.22)	5.29 (5.27, 5.30)	3.56 (3.54, 3.57)
UB^s	10.23 (10.19, 10.26)	7.69 (7.63, 7.74)	5.91 (5.86, 5.97)	10.26 (10.23, 10.28)	7.72 (7.68, 7.77)	5.95 (5.89, 6.00)
LB^m	9.51 (9.49, 9.53)	6.07 (6.05, 6.09)	4.05 (4.04, 4.06)	9.51 (9.50, 9.53)	6.08 (6.06, 6.09)	4.06 (4.05, 4.07)
UB^m	9.54 (9.51, 9.57)	6.20 (6.15, 6.26)	4.19 (4.12, 4.25)	9.55 (9.52, 9.57)	6.22 (6.17, 6.26)	4.20 (4.14, 4.26)
V^u	15.91	8.45	5.47	16.00	8.49	5.49

Table 3: **No short sales and no borrowing.** The four parameter sets are defined in Table 1. All results are computed for the initial value of the state variable $X_0 = 0$. The rows marked LB^s and LB^m report the estimates of the expected utility achieved by using the static and myopic portfolio strategies, respectively. The estimates are based on $1,000,000$ independent simulations. Expected utility is reported as a continuously compounded certainty equivalent return. Approximate 95% confidence intervals are reported in parentheses. The rows marked UB^s and UB^m report the analogous results for the upper bound on the true value function computed according to the procedure described in Section 4. The row marked V^u reports the optimal value function for the model with incomplete markets and without the portfolio constraints.

	$\gamma = 1.5$	$\gamma = 3$	$\gamma = 5$

<div align="center">Parameter set 1</div>

	$\gamma = 1.5$	$\gamma = 3$	$\gamma = 5$
LB^s	7.47	4.62	3.27
	(7.41, 7.53)	(4.59, 4.65)	(3.25, 3.29)
UB^s	10.35	6.28	4.32
	(10.21, 10.49)	(6.06, 6.50)	(4.02, 4.62)
LB^m	9.04	5.58	3.89
	(8.96, 9.12)	(5.53, 5.62)	(3.86, 3.92)
UB^m	9.14	5.77	4.04
	(9.02, 9.27)	(5.56, 5.98)	(3.75, 4.33)
V^u	9.44	5.95	4.19

<div align="center">Parameter set 2</div>

	$\gamma = 1.5$	$\gamma = 3$	$\gamma = 5$
LB^s	6.52	3.67	2.58
	(6.44, 6.60)	(3.63, 3.72)	(2.56, 2.61)
UB^s	10.76	5.85	3.88
	(10.60, 10.92)	(5.67, 6.03)	(3.69, 4.08)
LB^m	8.92	4.86	3.30
	(8.82, 9.01)	(4.81, 4.91)	(3.26, 3.33)
UB^m	9.01	4.98	3.38
	(8.87, 9.15)	(4.82, 5.14)	(3.20, 3.56)
V^u	9.02	4.92	3.33

<div align="center">Parameter set 3</div>

LB^s	11.19	6.76	4.61
	(11.12, 11.26)	(6.72, 6.80)	(4.59, 4.64)
UB^s	14.68	8.89	6.09
	(14.50, 14.85)	(8.61, 9.18)	(5.70, 6.48)
LB^m	12.80	7.82	5.34
	(12.71, 12.89)	(7.77, 7.88)	(5.30, 5.38)
UB^m	12.99	8.22	5.73
	(12.83, 13.14)	(7.95, 8.49)	(5.35, 6.11)
V^u	16.79	10.32	7.06

<div align="center">Parameter set 4</div>

LB^s	9.71	5.29	3.55
	(9.61, 9.81)	(5.23, 5.34)	(3.52, 3.59)
UB^s	13.67	7.32	4.75
	(13.49, 13.85)	(7.10, 7.54)	(4.50, 5.01)
LB^m	11.36	6.10	4.04
	(11.25, 11.47)	(6.04, 6.16)	(4.00, 4.08)
UB^m	11.47	6.21	4.09
	(11.30, 11.63)	(6.02, 6.41)	(3.86, 4.33)
V^u	15.91	8.45	5.47

Table 4: **No short sales.** The four parameter sets are defined in Table 1 and the problem horizon is $T = 5$ years. All results are computed for the initial value of the state variable $X_0 = 0$. The rows marked LB^s and LB^m report the estimates of the expected utility achieved by using the static and myopic portfolio strategies, respectively. The estimates are based on $100,000$ independent simulations. Expected utility is reported as a continuously compounded certainty equivalent return. Approximate 95% confidence intervals are reported in parentheses. The rows marked UB^s and UB^m report the analogous results for the upper bound on the true value function computed according to the procedure described in Section 4. The row marked V^u reports the optimal value function for the model with incomplete markets and without the portfolio constraints.