Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Tax-Aware Dynamic Asset Allocation

Martin Haugh, Garud Iyengar, Chun Wang

Operations Research
Articles in Advance, pp. 1–18
ISSN 0030-364X (print) | ISSN 1526-5463 (online)

http://dx.doi.org/10.1287/opre.2016.1517

Tax-Aware Dynamic Asset Allocation

Martin Haugh, Garud Iyengar, Chun Wang
Department of Industrial Engineering (IE) and Operations Research (OR), Columbia University, New York, New York 10027
{mh2078@columbia.edu, garud@ieor.columbia.edu, cw2519@columbia.edu}

We consider dynamic asset allocation problems where the agent is required to pay capital gains taxes on her investment gains. These are very challenging problems because the tax owed whenever a security is sold depends on the cost basis, and this results in high-dimensional problems, which cannot be solved exactly except in the case of very stylized problems with just one or two securities and relatively few time periods. In this paper, we focus on exact and average cost-basis problems, make the *limited use* of *losses* (LUL) assumption and develop simple heuristic trading policies for these problems when there are differential tax rates for long- and short-term gains and losses. We use information relaxation-based duality techniques to assess the performance of these trading policies by constructing unbiased lower and upper bounds on the (unknown) optimal value function. In numerical experiments with as many as 80 time periods and 25 securities we find our best suboptimal policy is within 3–10 basis points of optimality on a certainty equivalent (CE) annualized return basis. The principal contribution of this paper is in demonstrating that while the primal problem remains very challenging to solve exactly, we can easily solve very large dual problem instances. Moreover, dual tractability extends to standard problem variations, including problems with random time horizons, no wash sales constraints, intertemporal consumption and recursive utility, as well as the step-up feature of the U.S. tax code, among others.

Keywords: dynamic asset allocation; taxes; suboptimal control; duality; information relaxations.
Subject classifications: dynamic programming: applications; finance: portfolio, taxation.
Area of review: Financial Engineering.
History: Received March 2014; revisions received January 2016, March 2016; accepted April 2016. Published online in *Articles in Advance* June 24, 2016.

1. Introduction

Dynamic asset allocation problems have played a central problem in finance since the pioneering work of Samuelson (1969), Merton (1969) and Hakansson (1970). Since then, many researchers have studied the problem of how to dynamically allocate wealth among financial securities to optimize a given objective function, which is typically some combination of the expected utility of terminal wealth and lifetime consumption. Most of the work to date has considered problems with frictionless markets, and different problem formulations are typically obtained by varying the price dynamics and/or agent preferences. Relatively little work, however, has focused on asset allocation in the presence of capital gains taxes. This is not because the problem is unimportant. Indeed the problem of how to efficiently invest and rebalance a portfolio in the presence of capital gains taxes is of considerable interest to practitioners and academics alike. Rather, the problem has received relatively little attention because it is so challenging to solve. This is because the taxes that are owed whenever a security is sold generally depends on its *cost basis*, i.e., the price(s) at which the security was originally purchased. This feature results in high-dimensional and path-dependent problems, which can only be solved in extremely simple cases.

Constantinides (1984) was among the first to study the asset allocation problem with taxes. He showed that the optimal investment policy can be separated from the tax timing problem if short sales are allowed and are costless. In particular, when an agent needs to reduce her position in a stock with embedded capital gains, she prefers to short sell the stock rather than selling from her current holdings.[1] She therefore succeeds in rebalancing her portfolio without triggering a tax liability. In practice, of course, short sales incur collateral costs and may also not be permitted by the tax authorities.

Dybvig and Koo (1996) studied the asset allocation problem with taxes when short sales constraints are imposed so that the separation result of Constantinides does not apply. They formulated the problem using the so-called *exact* cost basis and solved a problem with just one risky stock and four time periods. More recently, DeMiguel and Uppal (2005) used a stochastic programming approach to solve problems with just one stock and 10 time periods, as well as problems with two stocks and 7 time periods. Unfortunately, solving larger problems is numerically intractable because the number of state variables and number of constraints grows exponentially with the number of time periods.

An alternative modeling approach is to use the *average* cost basis when determining tax liabilities. The average cost basis for a given security is the weighted average purchase price of the current holdings of the security in the portfolio. The current U.S. tax code requires agents to use exact cost basis for individual stocks, and gives agents an additional option to use the average cost basis for mutual fund shares. While using the average cost basis is generally suboptimal from the agent's point of view, it requires less record keeping, and from a modeling point of view, does not result in the exponential explosion of state variables that occurs when the exact cost basis is used. Average cost-basis problems are therefore more amenable to dynamic programming (DP) but are nonetheless still challenging to solve. In particular, due to the so-called curse of dimensionality, it is only possible to exactly solve problems with just one or two stocks even when the average cost basis is used.

Dammon et al. (2001, 2004) were among the first to consider the average cost-basis problem. They considered problems with just one stock and multiple time periods. Dammon et al. (2002) and Garlappi et al. (2001) considered the case of two stocks. Gallmeyer et al. (2006) solved a two stock problem as well as a problem with just one stock and a put option on that stock. It is worth noting that DeMiguel and Uppal (2005) showed that the certainty equivalent loss in wealth was small when using the average cost basis instead of the exact cost basis. Although they only showed this for the relatively small problem sizes they considered, one would expect this observation to hold more generally. Tahar et al. (2010) formulated the continuous time version of the model in Dammon et al. (2001) and studied properties of the value function in a one stock infinite-horizon problem.

We also note that most of the literature on tax-aware asset allocation assumes the so-called *full use of losses* (FUL) model where the net capital losses in any given period result in a tax rebate, which can be invested immediately. This is in contrast to the *limited use of losses* (LUL) model where the net capital losses do not result in an immediate tax rebate; instead, the net losses can only be used to offset capital gains in future periods. Only the LUL model is consistent with the tax codes encountered in practice but LUL problems are more challenging to solve than FUL problems. This is because in the LUL case, we need to keep track of the sum of prior losses that have not already been used to offset gains. Gallmeyer and Srivastava (2010) discussed the impact of the LUL model, while more recently, Ehling et al. (2010) solved a two-stock average cost-basis LUL problem with 80 time periods.[2]

In this paper, we focus on the exact and average cost-basis LUL cases since these problems are the most realistic and challenging to solve. We develop suboptimal and simple heuristic trading policies for these problems when there are differential tax rates for short- and long-term gains or losses. We then use duality techniques based on information relaxations to assess the performance of these trading policies by constructing unbiased lower and upper bounds on the (unknown) optimal value function. In numerical experiments with as many as 80 time periods and 25 securities, we find our best suboptimal policy is within 3–10 basis points of optimality on a certainty-equivalent annualized return basis. So instead of trying to solve these problems exactly, we settle for *provably* good suboptimal solutions.

The principal contribution of the paper is to demonstrate that much larger problems can now be tackled through the use of sophisticated optimization techniques and duality methods based on information relaxations (Brown et al. 2010, Rogers 2007). To the best of our knowledge, we are the first to construct valid dual bounds for tax-aware dynamic asset allocation problems. We show that the dual formulations of the exact cost-basis problems are much easier to solve than the corresponding primal problems. In contrast, while the average cost-basis (primal) problem is relatively easier to solve than the exact cost-basis (primal) problem, the dual of the average cost-basis problem is nonconcave and thus more difficult to solve. We use the polyhedral branch-and-cut approach of Tawarmalani and Sahinidis (2005) to bound the dual problem instances so that valid lower and upper bounds for the average cost-basis problem can still be obtained. This approach should be useful in other applications where the dual problem instances can sometimes fail to be convex. Moreover, we show the tractability of the dual problem instances extends to standard problem variations, including problems with random time horizons, intertemporal consumption with recursive utility, no wash-sales constraints and the step-up feature of the U.S. tax code, among others.

In addition, our numerical results support the conclusion of DeMiguel and Uppal (2005). In particular, we show numerically that the certainty equivalent loss in wealth is very small when using the average cost basis instead of the exact cost basis. Finally, we note that our main numerical application includes differential tax rates for short- and long-term capital gains or losses. In addition to being more realistic, allowing for differential tax rates allows us to evaluate simple heuristic policies that are employed in practice. For example, a very common and sensible rule of thumb is that one should always avoid paying short-term capital gains taxes. We can easily evaluate simple heuristic policies containing such rules in our differential tax rate framework. Moreover, in our numerical results, we find that such heuristics can perform very well.

The remainder of this paper is organized as follows. In Section 2, we formulate the tax-aware dynamic asset allocation problem, focusing mainly on the exact cost-basis LUL problem with differential tax rates on short- and long-term gains. We describe a suboptimal policy and some simple heuristics for tackling this problem in Section 3. In Section 4, we describe how we use recently developed techniques based on information relaxations to construct dual bounds for these tax-aware portfolio optimization problems. Numerical results are presented in Section 5. We consider the average cost-basis problem in Section 6 and conclude in Section 7. Appendix A provide further solution details as well as the no-tax problem formulation while various extensions are considered in an online appendix (available as supplemental material at http://dx.doi.org/10.1287/opre.2016.1517).

2. Main Problem Formulation

Our basic framework is an extension of the multiple assets model of Dybvig and Koo (1996) and DeMiguel and Uppal (2005). We consider the problem of a risk-averse agent who can trade in multiple assets over a finite-time horizon. We assume that borrowing and short sales are prohibited so that the aforementioned separation result of Constantinides (1984) does not apply. The agent must pay taxes on any realized capital gains that she incurs and the tax rate will, in general, depend on whether the gains are short or long term. We define short-term gains (and losses) to be those realized from selling securities that have been held for m_{st} periods or less, where $0 < m_{\mathrm{st}} < T$ and the subscript "st" denotes "short term." In contrast, long-term gains and losses are those realized from selling securities that have been held for more than m_{st} periods. Let τ^s and τ^l denote the short- and long-term tax rates, respectively. We assume $0 < \tau^l < \tau^s$, which, of course, is typically true in practice as governments seek to encourage long-term investing.

We assume a discrete-time economy with equally spaced trading dates indexed by $t = 0, 1, \ldots, T$. The market consists of a risk-free cash account that earns an after -tax risk-free return of r_0 per period, and K risky securities. Let $p_{t,k}$ denote the price of the kth risky security at time t, and let $\mathbf{p}_t := [p_{t,1}, \ldots, p_{t,K}]'$ denote the vector of security prices at time t. Without loss of generality, the initial security prices are set to one, i.e., $\mathbf{p}_0 = [1, \ldots, 1]'$. Let $r_{t,k}$ denote the stochastic gross return of the kth risky security between times $t-1$ and t, and let $\mathbf{r}_t := [r_{t,1}, \ldots, r_{t,K}]'$ denote the vector of gross returns. Then,

$$\mathbf{p}_t = \mathbf{p}_{t-1} \cdot \mathbf{r}_t, \tag{1}$$

where $\mathbf{p}_{t-1} \cdot \mathbf{r}_t$ is the componentwise multiplication of $\mathbf{p}_{t-1}$ and $\mathbf{r}_t$. We allow for the possibility that $\mathbf{r}_{t+1}$ is driven by a Markovian state vector $\mathbf{z}_t$, so that $\mathbf{r}_{t+1}$ is independent of $\mathbf{r}_j$ for all $j \leqslant t$ conditional on $\mathbf{z}_t$.

Let b_t denote the cash account holding at time t, and let $n_{j,t,k}$ denote the number of units of the kth security that were purchased at time $j \leqslant t$ and still held in the agent's portfolio *after* trading at time t for $t = 0, \ldots, T$. The *exact* cost basis of these $n_{j,t,k}$ units is then given by $p_{j,k}$. Let $\mathbf{n}_{j,t} := [n_{j,t,1}, \ldots, n_{j,t,K}]'$, $0 \leqslant j \leqslant t$ denote the stock holdings at time t. Since we do not allow short sales or borrowing, we must have that

$$\begin{aligned} &b_t \geqslant 0, \quad \mathbf{n}_{t,t} \geqslant \mathbf{0}, \quad t \geqslant 0, \quad \text{and} \\ &\mathbf{n}_{j,t-1} \geqslant \mathbf{n}_{j,t} \geqslant \mathbf{0}, \quad t \geqslant 1, \quad 0 \leqslant j < t. \end{aligned} \tag{2}$$

Note that the sequence of vectors $\{(b_t, \mathbf{n}_{j,t})\colon 0 \leqslant j \leqslant t\}_{t=0}^T$ implicitly defines a trading strategy. Next, we define the tax implications of this trading strategy.

It is standard in the tax-aware portfolio selection literature to assume FUL whereby realized capital losses earn an immediate tax rebate. See, for example, Dybvig and Koo (1996), Dammon et al. (2001, 2002, 2004), Garlappi et al. (2001), DeMiguel and Uppal (2005), Gallmeyer et al. (2006), and Tahar et al. (2010). Furthermore, these papers typically ignore how long the securities have been held and assume that all capital gains are taxed at the same rate. In practice, however, tax codes only allow for the LUL in that an agent can only use realized capital losses to offset realized capital gains. Any unused capital losses can then be carried forward to offset future capital gains. Moreover, tax codes often apply differential tax rates to short- and long-term capital gains. We will focus here on the more challenging problem, which assumes LUL and differential tax rates for short- and long-term gains. In addition, the current U.S. tax code requires that short- and long-term losses must be used first to offset gains of the same type. If short-term losses exceed short-term gains, the agent is allowed[3] to use the excess short-term losses to offset excess long-term gains.

We model these tax rules as follows. Let l_{t-1}^s (resp. l_{t-1}^l) denote the accumulated unused short-term (resp. long-term) losses carried forward from time $t-1$. We let

$$c_t^s = \sum_{j=t-m_{\mathrm{st}}}^{t-1} (\mathbf{p}_t - \mathbf{p}_j)'(\mathbf{n}_{j,t-1} - \mathbf{n}_{j,t}) \quad \text{and}$$

$$c_t^l = \sum_{j=0}^{t-m_{\mathrm{st}}-1} (\mathbf{p}_t - \mathbf{p}_j)'(\mathbf{n}_{j,t-1} - \mathbf{n}_{j,t})$$

denote the short-term and long-term net realized proceeds, respectively, at time t. (Note that $\mathbf{n}_{j,t-1} - \mathbf{n}_{j,t}$ are the units of securities with cost-basis $\mathbf{p}_j$ that are sold at time t.) The short-term gains g_t^s and the excess "intermediate"[4] short-term losses $\hat{l}_t^s$ at time t are then given by

$$g_t^s = \max\{c_t^s + l_{t-1}^s, 0\}, \quad \hat{l}_t^s = \min\{c_t^s + l_{t-1}^s, 0\}. \tag{3}$$

Prepaying additional taxes early is never optimal in this model, since one can invest that amount in the risk-free cash account, which earns a nonnegative interest rate r_0. Furthermore, since $\tau^s > \tau^l$, it is never optimal to use short-term losses to offset long-term gains before completely offsetting short-term gains. Therefore the optimal solution will never allow for $g_t^s > 0$ and $\hat{l}_t^s < 0$ to occur simultaneously, and thus (3) can be linearized as follows:

$$g_t^s + \hat{l}_t^s = c_t^s + l_{t-1}^s = \sum_{j=t-m_{\mathrm{st}}}^{t-1} (\mathbf{p}_t - \mathbf{p}_j)'(\mathbf{n}_{j,t-1} - \mathbf{n}_{j,t}) + l_{t-1}^s.$$

Next, we consider the long-term gains and losses. The unused long-term losses l_t^l and the "intermediate" long-term gains $\hat{g}_t^l$ at time t are given by

$$\hat{g}_t^l = \max\{c_t^l + l_{t-1}^l, 0\}, \quad l_t^l = \min\{c_t^l + l_{t-1}^l, 0\}. \tag{4}$$

Again, because prepaying additional taxes early is suboptimal, it will never be optimal to have $\hat{g}_t^l > 0$ and $l_t^l < 0$ to occur simultaneously. Thus (4) can be replaced by the linear constraint

$$\hat{g}_t^l + l_t^l = c_t^l + l_{t-1}^l = \sum_{j=0}^{t-m_{\mathrm{st}}-1} (\mathbf{p}_t - \mathbf{p}_j)'(\mathbf{n}_{j,t-1} - \mathbf{n}_{j,t}) + l_{t-1}^l.$$

The linearization of the dynamics of g_t^s and $\hat{g}_t^l$ is crucial to ensure that the resulting portfolio selection problem is convex and therefore tractable.

At this point, we have used only short-term (resp. long-term) losses to offset short-term (resp. long-term) gains. We now allow[5] excess short-term losses to offset excess long-term gains via the constraints

$$g_t^l + l_t^s = \hat{g}_t^l + \hat{l}_t^s, \quad 0 \leqslant g_t^l \leqslant \hat{g}_t^l, \quad \hat{l}_t^s \leqslant l_t^s \leqslant 0,$$

where the two inequalities above guarantee that g_t^l and l_t^s move in the correct direction. Note that we have implicitly assumed that unused losses l_t^s and l_t^l can be carried forward indefinitely and don't have a "use by date." Since $\tau^s > \tau^l$, the agent may prefer rolling over the excess short-term losses rather than fully using them to offset long-term gains; the constraint above allows for this possibility.

The boundary conditions on $(g_t^s, \hat{l}_t^s, \hat{g}_t^l, l_t^l)$ are

$$\begin{aligned} &g_t^s \geqslant 0, \quad \hat{l}_t^s \leqslant 0, \quad t \geqslant 1, \quad \text{and} \\ &\hat{g}_t^l \geqslant 0, \quad l_t^l \leqslant 0, \quad t > m_{\text{st}}. \end{aligned} \tag{5}$$

We also need the following initial conditions:

$$l_t^s = \hat{l}_t^s, \quad g_t^l = l_t^l = 0, \quad 1 \leqslant t \leqslant m_{\text{st}}, \quad \text{and} \quad l_0^s = 0. \tag{6}$$

They reflect the fact that there is no long-term gains from trading in the first m_{st} periods and our assumption that the initial short- and long-term carried losses are zero. This, of course, is easy to change if we begin at time $t = 0$ with some short- and/or long-term carried losses.

The agent is assumed to have an initial wealth w_0, and the objective is to maximize the agent's expected utility of terminal after-tax wealth b. Note that under this objective, we require all risky security positions to be liquidated at the end of the time horizon, T, and any taxes owed at this time must be paid. We assume here that the agent's utility function U belongs to the CRRA class of utility functions so that

$$U(b) := \begin{cases} \dfrac{b^{1-\gamma}}{1-\gamma}, & \gamma > 0 \text{ and } \gamma \neq 1, \\ \ln(b), & \gamma = 1, \end{cases} \tag{7}$$

where γ is the coefficient of relative risk aversion. Thus the tax-aware asset allocation problem can then be formulated as

$$\max_{\substack{(b_t, \mathbf{n}_{j,t}, g_t^s, l_t^s, g_t^l, l_t^l, \hat{g}_t^l, \hat{l}_t^s) \in \mathcal{F}_t, \\ t=0,\dots,T}} \mathbb{E}_0\left[\frac{b_T^{1-\gamma}}{1-\gamma}\right] \tag{8}$$

s.t.

$$b_0 + \mathbf{p}_0' \mathbf{n}_{0,0} = w_0, \tag{9}$$

$$b_t + \sum_{j=0}^{t} \mathbf{p}_t' \mathbf{n}_{j,t} + \tau^s g_t^s + \tau^l g_t^l = b_{t-1} r_0 + \sum_{j=0}^{t-1} \mathbf{p}_t' \mathbf{n}_{j,t-1}, \quad t \geqslant 1, \tag{10}$$

$$g_t^s + \hat{l}_t^s = \sum_{j=t-m_{\text{st}}}^{t-1} (\mathbf{p}_t - \mathbf{p}_j)'(\mathbf{n}_{j,t-1} - \mathbf{n}_{j,t}) + l_{t-1}^s, \quad t \geqslant 1, \tag{11}$$

$$\hat{g}_t^l + l_t^l = \sum_{j=0}^{t-m_{\text{st}}-1} (\mathbf{p}_t - \mathbf{p}_j)'(\mathbf{n}_{j,t-1} - \mathbf{n}_{j,t}) + l_{t-1}^l, \quad t > m_{\text{st}}, \tag{12}$$

$$g_t^l + l_t^s = \hat{g}_t^l + \hat{l}_t^s \quad t > m_{\text{st}}, \tag{13}$$

$$0 \leqslant g_t^l \leqslant \hat{g}_t^l \quad t > m_{\text{st}}, \tag{14}$$

$$\hat{l}_t^s \leqslant l_t^s \leqslant 0 \quad t > m_{\text{st}}, \tag{15}$$

together with the nonnegativity trade constraints (2), the boundary and initial conditions on the tax variables (5) and (6), as well as security price and state variable dynamics. Note that we use $\{\mathcal{F}_t\}_{t=0,\dots,T}$ to denote the filtration generated by the security price vectors as well as any other state variables in the model. (9) is the initial budget constraint, and budget constraints for $t \geqslant 1$ are given by (10) where the right-hand side is the pretrade wealth at time t and the sum of the first two terms on the left-hand side represents the posttrade wealth after paying taxes of $\tau^s g_t^s$ and $\tau^l g_t^l$ on the short- and long-term gains, respectively. Since the short- and long-term capital gains g_t^s and g_t^l are constrained to be nonnegative, the agent can never receive a tax rebate in (10). The decision variables $\mathbf{n}_{j,t}$ allow the agent to track the exact cost basis of each share in the portfolio. Recall that short selling and borrowing via the cash account are ruled out by nonnegativity constraints (2). The constraints (11)–(15), the boundary conditions (5), and the initial conditions (6) define the dynamics of the short- and long-term gains. Since the objective is concave, and all constraints are linear, (8) is a convex optimization problem. Note also that in this formulation, we allow[6] so-called *wash sales* so that securities can be sold and purchased back in the same period. In the online appendix, we describe how wash sales can be easily handled when solving dual problem instances.

To simplify later problem formulations we define $\mathbf{x}_t := [b_t \; \mathbf{n}_{0,t}' \; \dots \; \mathbf{n}_{t,t}' \; g_t^s \; l_t^s \; g_t^l \; l_t^l \; \hat{g}_t^l \; \hat{l}_t^s]'$ and $\mathbf{p}_{0:t} := [\mathbf{p}_0' \; \dots \; \mathbf{p}_t']'$, the vector of all risky security prices up to and including time t. Together $\mathbf{x}_{t-1}$ and $\mathbf{p}_{0:t}$ completely describe the positions and cost basis of the agent's portfolio just before trading at time t. We denote the set of feasible trades at time t by $\mathbb{X}_t$. Thus

$$\mathbb{X}_0 := \{\mathbf{x}_0 \mid \mathbf{x}_0 \in \mathcal{F}_0 \text{ satisfies (2) at } t = 0 \text{ and (9)}\} \tag{16}$$

$$\mathbb{X}_t(\mathbf{x}_{t-1}, \mathbf{p}_{0:t}) := \{\mathbf{x}_t \mid \mathbf{x}_t \in \mathcal{F}_t \text{ satisfies (2), (5), (6), and (10)–(15)}\}, \quad t \geqslant 1. \tag{17}$$

Unfortunately, it is impossible to solve the exact cost-basis LUL problem due to the large number of state variables, constraints, and path dependence induced by the need to keep track of the cost basis for each security. Instead, we seek good suboptimal policies, and computing them is the subject of Section 3.

3. Feasible Investment Policies

We now discuss several feasible policies for the tax-aware asset allocation problem that we formulated in Section 2. We note that each of these policies can be easily simulated so that it is straightforward to obtain unbiased estimators of their performances. Given that these policies are feasible, it is also clear that these estimators are lower or *primal* bounds on the optimal value function, V_0^{LUL}. We begin with policies based on the solution to the no-tax problem. The no-tax problem is formulated in Appendix A.1 and is relatively easy to solve numerically as long as the state vector $\mathbf{z}_t$ is low dimensional. When discussing no-tax portfolios, we will use $\tilde{b}_t$ and $\tilde{\mathbf{n}}_t$ to denote the posttrade *fractions* of wealth w_t invested in the cash account and risky securities, respectively. This is in contrast to the formulation in Section 2, where b_t is the number of dollars invested in the cash account and $\mathbf{n}_{j,t}$ denotes the number of *units* of the risky securities purchased at time j and still held after trading at time t.

3.1. The Tax Blind Policy

Under the tax blind policy, we assume the agent simply completely ignores capital gains taxes. In particular, at each time period, she liquidates her entire portfolio of risky securities, pays short-term capital gains taxes on her net gains, and then trades to the no-tax optimal portfolio weights $(\tilde{b}_t^*, \tilde{\mathbf{n}}_t^*)$. This strategy is obviously very naive but it does serve as a baseline policy that we can use to compare the performance of our other strategies.

3.2. The Tax-Aware Heuristic Policy

The tax-aware policy is a heuristic policy where the agent follows two basic principles:

(a) Harvest all losses in every time period.

(b) Always avoid paying short-term capital gains taxes.

Subject to these two rules, the agent attempts to trade toward the no-tax optimal portfolio weights. More specifically, in each time period, the heuristic takes the following steps:

(1) Harvest all short- and long-term losses by selling any shares whose cost basis is higher than the current price. Let l^s (resp. l^l) denote the sum of the harvested short-term (resp. long-term) losses and unused short-term (resp. long-term) losses carried from previous periods.

(2) Let w_t denote the pretrade portfolio wealth, and recall that $(\tilde{b}_t^*, \tilde{\mathbf{n}}_t^*)$ are the optimal portfolio weights implied by the no-tax problem. Then, the target dollar position for the cash account and securities are given by $w_t\tilde{b}_t^*$ and $w_t\tilde{\mathbf{n}}_t^*$, respectively.

(3) Sell shares of all risky securities with dollar position higher than the target dollar position. Begin by first selling shares with embedded long-term gains. If there are lots with more than one cost basis, sell in proportion to the dollar amount held in each cost basis.

At this point, any shares that still need to be sold to reach the target dollar position have embedded short-term gains. Let $\hat{s}_k$ denote the total dollar amount of shares with short-term gains that need to be sold for security $k = 1, \ldots, K$. Sell $s_k = \min\{(\hat{s}_k / \sum_j \hat{s}_j) l^s, \hat{s}_k\}$ dollars worth of shares of security k.

(4) The short-term gains realized in the previous step can be fully offset by l^s, and thus there is no need to pay short-term capital gains taxes.

Offset the realized long-term gains by first using l^l and then by using any remaining unused short-term losses. Pay long-term taxes if the net long-term gains are positive.

(5) If the position in the cash account is higher than the target dollar amount, designate the difference $\hat{b}_t$ as free capital available for investing in the risky securities. Let $\hat{d}_k$ denote the dollar amount that needs to be purchased to reach the target dollar position for security $k = 1, \ldots, K$. Purchase[7] $d_k = (\hat{d}_k / \sum_k \hat{d}_k)\hat{b}_t$ dollars worth of security k.

To better understand[8] the mechanics of this policy, we provide some examples in Appendix A.2.

3.3. The Rolling Buy-and-Hold Policy

The rolling buy-and-hold (RBH) policy is constructed as follows. At each time t, the agent assumes she can trade at the current time t, but that she will next be allowed to trade again *only* at time T when her portfolio is liquidated. Suppose the agent's holdings after trade at time t are given by $(b_t, \mathbf{n}_{0,t}, \ldots, \mathbf{n}_{t,t})$. Then, the realized cash position b_T at time T is given by

$$b_T = b_t r_0^{T-t} + \sum_{j=0}^{t} \mathbf{p}_T' \mathbf{n}_{j,t} - \tau^s g_T^s - \tau^l g_T^l,$$

where g_T^s is the realized short-term gains and g_T^l is the realized long-term gains. When the current period $t < T - m_{\text{st}}$, all risky positions $\mathbf{n}_{j,t}$ will become long-term holdings by time T, and $g_T^s = 0$. Thus, short-term gains g_T^s are possibly nonzero for only $t \in \{T - m_{\text{st}}, \ldots, T\}$. Since m_{st} is typically small,[9] in practice, we choose to combine short- and long-term gains and pay taxes at the long-term tax rate in the final m_{st} periods, and thus we solve the following simplified version of the problem to select the holdings $(b_t, \mathbf{n}_{0,t}, \ldots, \mathbf{n}_{t,t})$:

$$\max_{\mathbf{x}_t, b_T, g_T} \; \mathbb{E}_t\left[\frac{b_T^{1-\gamma}}{1-\gamma}\right] \tag{18}$$

$$\text{s.t.} \quad \mathbf{x}_t \in \mathbb{X}_t(\mathbf{x}_{t-1}, \mathbf{p}_{0:t}),$$

$$b_T = b_t r_0^{T-t} + \sum_{j=0}^{t} \mathbf{p}_T' \mathbf{n}_{j,t} - \tau^l g_T, \tag{19}$$

$$g_T \geqslant \sum_{j=0}^{t} (\mathbf{p}_T - \mathbf{p}_j)' \mathbf{n}_{j,t} + l_t^l + l_t^s, \tag{20}$$

$$g_T \geqslant 0, \tag{21}$$

where $\mathbb{X}_t(\mathbf{x}_{t-1}, \mathbf{p}_{0:t})$ denotes the set of feasible trades at time t defined in (16) and (17). The constraint (19) defines the terminal wealth b_T, which assumes no trading is allowed between times t and T, and all taxable gains are long-term gains. Note that we do not, however, violate the LUL rule with differential short- and long-term tax rates when actually implementing the strategy. In particular, the trading decisions at time t are constrained in the feasible set $\mathbb{X}_t$, where short- and long-term taxes are handled correctly.

After the agent implements the optimal solution of (18)–(21), she holds this position until time $t+1$ at which point she will rebalance the portfolio by solving (18)–(21) again but now with $t=t+1$.

3.3.1. Solution Approach. We obtain an unbiased estimate for the value function of the RBH policy by averaging the performance over I sample paths. Over each path i and each time instant t along the path, we need to compute the optimal trade by solving (18). The expectation in the objective function of (18) cannot be evaluated analytically, so instead, we approximate it by taking samples. In particular, we use low-discrepancy sequences[10] (LDS) to generate M scenarios $\mathbf{p}_T^{(m)}$, $m=1,\ldots,M$ for the random time T security prices. At time t, we choose the RBH trades by solving the following approximation to (18):

$$\max_{\mathbf{x}_t,\, b_T^{(m)},\, g_T^{(m)}} \quad \frac{1}{M}\sum_{m=1}^{M}\frac{(b_T^{(m)})^{1-\gamma}}{1-\gamma} \tag{22}$$

$$\text{s.t.}\quad \mathbf{x}_t \in \mathbb{X}_t(\mathbf{x}_{t-1}, \mathbf{p}_{0:t}),$$

$$b_T^{(m)} = b_t r_0^{T-t} + \sum_{j=0}^{t}\mathbf{p}_T^{(m)\prime}\mathbf{n}_{j,t} - \tau^l g_T^{(m)}, \quad 1 \leqslant m \leqslant M, \tag{23}$$

$$g_T^{(m)} \geqslant \sum_{j=0}^{t}(\mathbf{p}_T^{(m)} - \mathbf{p}_j)'\mathbf{n}_{j,t} + l_t^l + l_t^s, \quad 1 \leqslant m \leqslant M, \tag{24}$$

$$g_T^{(m)} \geqslant 0 \quad 1 \leqslant m \leqslant M, \tag{25}$$

where $\tau^l g_T^{(m)}$ is the capital gains taxes paid after liquidating the portfolio at time T in the mth scenario and $b_T^{(m)}$ is the resulting after-tax cash holding. The optimization problem (22)–(25) has[11] $2M+K(t+1)+7$ decision variables and $2M+6$ constraints. Since the objective function (22) is concave and all constraints are linear, (22) is a convex optimization problem, and in principle, can be solved using standard[12] solvers.

We need the number of scenarios M for the time T security prices to be large to adequately approximate the expectation in (18). For the numerical results of Section 5 where we had 25 risky securities and 80 time periods, we chose $M=10{,}000$ scenarios. MOSEK then required 15 minutes[13] to solve (22)–(25) for all trading periods on a given sample path. If one uses $N=5{,}000$ sample paths to estimate the value of the RBH policy, the total running time is then approximately 52 days! We therefore need an alternative approach to solving the approximate RBH problem.

Toward this end, recall that (24) and (25) are a reformulation of the LUL tax rule

$$g_T^{(m)}(\mathbf{x}_t) = \max\left\{\sum_{j=0}^{t}(\mathbf{p}_T^{(m)} - \mathbf{p}_j)'\mathbf{n}_{j,t} + l_t^l + l_t^s, 0\right\}. \tag{26}$$

Thus (22) can be reformulated as

$$\max_{\mathbf{x}_t} \quad \frac{1}{M(1-\gamma)} \cdot \sum_{m=1}^{M}\left(b_t r_0^{T-t} + \sum_{j=0}^{t}\mathbf{p}_T^{(m)\prime}\mathbf{n}_{j,t} - \tau^l g_T^{(m)}(\mathbf{x}_t)\right)^{1-\gamma} \tag{27}$$

$$\text{s.t.}\quad \mathbf{x}_t \in \mathbb{X}_t(\mathbf{x}_{t-1}, \mathbf{p}_{0:t}),$$

where there are now a total of $K(t+1)+7$ variables and only 6 constraints, namely, (10) and (11)–(15), which are implicitly present in $\mathbb{X}_t(\mathbf{x}_{t-1}, \mathbf{p}_{0:t})$. Note that the size of the problem (27) is now independent of M and much smaller than that of problem (22)–(25). However, the function $g_T^{(m)}(\mathbf{x}_t)$ is nonsmooth, and therefore standard solvers that rely on first- and second-order derivatives cannot be used to solve (27).

The nonlinear nonsmooth function $\max\{a,0\}$ can be approximated by the smooth function

$$\phi_\theta(a) = (1/\theta)\ln(e^{\theta a}+1),$$

$\theta > 0$, in the sense that $\max\{a,0\} \leqslant \phi_\theta(a) \leqslant \max\{a,0\} + 2/\theta$. Thus $\phi_\theta(a)$ converges uniformly to $\max\{a,0\}$ as $\theta \to \infty$. We therefore approximate $g_T^{(m)}(\mathbf{x})$ in (27) by the smooth function

$$\tilde{g}_T^{(m)}(\mathbf{x}_t) = \phi_\theta(g_T^{(m)}(\mathbf{x})) = \frac{1}{\theta}\ln\left(\exp\left(\theta\left(\sum_{j=0}^{t}(\mathbf{p}_T^{(m)} - \mathbf{p}_j)'\mathbf{n}_{j,t} + l_t^l + l_t^s\right)\right)+1\right) \tag{28}$$

to obtain the following smooth approximate optimization problem:

$$\max_{\mathbf{x}_t} \quad \frac{1}{M(1-\gamma)} \cdot \sum_{m=1}^{M}\left(b_t r_0^{T-t} + \sum_{j=0}^{t}\mathbf{p}_T^{(m)\prime}\mathbf{n}_{j,t} - \tau^l \tilde{g}_T^{(m)}(\mathbf{x}_t)\right)^{1-\gamma} \tag{29}$$

$$\text{s.t.}\quad \mathbf{x}_t \in \mathbb{X}_t(\mathbf{x}_{t-1}, \mathbf{p}_{0:t}).$$

In our numerical experiments, we solved (29) with θ set equal to 50 using a sequential quadratic programming (SQP) approach, the details of which can be found in Appendix A.4. We typically found that our SQP algorithm converged after only two or three iterations. The time required to solve the resulting quadratic programming problems was significantly smaller than that required to solve the general convex programming problem (22). Returning to our earlier example of 25 stocks and 80 time periods with $M = 10{,}000$ scenarios, the SQP approach requires approximately 20 seconds to solve all 80 problems on one sample path.[14] Computing the RBH policy along 5,000 sample paths to estimate the value function therefore takes approximately 28 hours. This corresponds to approximately a 98% improvement in the solution time.

4. Evaluating Suboptimal Policies via Information Relaxations

The performance of any feasible policy for the exact cost-basis LUL problem in (8) is clearly a lower bound on the optimal value function V_0^{LUL}. However, since V_0^{LUL} is not computable, we cannot assess the quality of any feasible policy. We can overcome this by computing valid upper bounds for the value function. We show in Appendix A.1 that the no-tax value function V_0^N is an upper bound for V_0^{LUL}. This bound, however, is typically very weak, and therefore provides relatively little information regarding the quality of the suboptimal feasible policies. In this section, we show how to construct tighter upper (or dual) bounds using the recently developed information relaxation approach of Rogers (2007) and, in particular, Brown et al. (2010). We refer the reader to Brown et al. (2010) for the theoretical details underpinning the methods described in this section.

4.1. Dual Problem Formulation

Let $\mathbf{p} := [\mathbf{p}_0' \dots \mathbf{p}_T']'$, $\mathbf{z} := [\mathbf{z}_0' \dots \mathbf{z}_T']'$, and $\mathbf{x} := [\mathbf{x}_0' \dots \mathbf{x}_T']'$ denote the entire sequence of security prices, market states, and trade decision vectors, respectively. A trading policy can be interpreted as a function $\mathbf{x}(\mathbf{p}, \mathbf{z})$ that maps sequences of prices $\mathbf{p}$ and market states $\mathbf{z}$ to a sequence of trading decisions $\mathbf{x}$. A feasible trading policy is one where each individual decision $\mathbf{x}_t$ is in the set of feasible trades, $\mathbb{X}_t$. A feasible adapted policy is a feasible policy that is $\mathcal{F}_t$-adapted. Let $\mathcal{X}^L$ denote the set of all feasible adapted policies for the exact cost-basis LUL problem. The agent wants to compute a feasible adapted policy that maximizes the expected utility of the after-tax terminal wealth. Thus the optimization problem (8) can be reformulated as

$$V_0^{\mathrm{LUL}} = \max_{\mathbf{x} \in \mathcal{X}^L} \left\{ \mathbb{E}_0\left[\frac{b_T^{1-\gamma}}{1-\gamma} \right] \right\}. \tag{30}$$

We will call a function $\pi(\mathbf{x})$ a *dual feasible penalty* if $\mathbb{E}_0[\pi(\mathbf{x}(\mathbf{p},\mathbf{z}))] \leqslant 0$ for all feasible $\mathcal{F}_t$-adapted policies, $\mathbf{x}(\mathbf{p},\mathbf{z}) \in \mathcal{X}^L$. It is then clear that

$$V_0^{\mathrm{LUL}} \leqslant \max_{\mathbf{x} \in \mathcal{X}^L} \left\{ \mathbb{E}_0\left[\frac{b_T^{1-\gamma}}{1-\gamma} - \pi(\mathbf{x}) \right] \right\}.$$

Suppose at time $t = 0$ the agent has *perfect information* of all future security prices and market states before she makes any trading decisions. It then follows that

$$\begin{aligned} V_0^{\mathrm{LUL}} &\leqslant \max_{\mathbf{x} \in \mathcal{X}^L} \left\{ \mathbb{E}_0\left[\frac{b_T^{1-\gamma}}{1-\gamma} - \pi(\mathbf{x}) \right] \right\} \\ &\leqslant \mathbb{E}_0\left[\max_{\mathbf{x} \in \mathbb{X}} \left\{ \frac{b_T^{1-\gamma}}{1-\gamma} - \pi(\mathbf{x}) r \right\} \right] =: V_{\mathrm{up}} \end{aligned} \tag{31}$$

where $\mathbb{X}$ denotes the set of feasible trades under the assumption of perfect information. Brown et al. (2010) and Rogers (2007) establish that strong duality holds in (31), i.e., there exist dual feasible penalties for which $V_{\mathrm{up}} = V_0^{\mathrm{LUL}}$. We will construct our dual penalty $\pi(\mathbf{x})$ using the solution to the no-tax asset allocation problem. We note at this point that the penalty will be linear in the actions $\mathbf{x}$. Moreover, it follows from results in Brown and Smith (2011) that the resulting dual bound will be at least as good as the upper bound provided by the no-tax solution. Brown and Haugh (2013) also provide some insight on how much of an improvement over the no-tax solution might be expected. Further details on these observations as well as the formulation and solution of the no-tax problem can be found in Appendices A.1 and A.3.

The dual bound of (31) leads itself to estimation via Monte Carlo simulation: we simply simulate I sample paths of the security prices and market states, and on each path, we solve the maximization problem inside the expectation in (31). If $V_{\mathrm{up}}^{(i)}$ is the optimal solution of the problem on the ith path, then $\sum_{i=1}^{I} V_{\mathrm{up}}^{(i)}/I$ is an unbiased estimate of V_{up}. Moreover, we can estimate V_{up} to any desired accuracy by choosing a large enough number of paths, I.

Suppose now we have simulated one sample path for the security prices and state variables. The inner maximization problem in (31) is a deterministic optimization problem. Since $b^{1-\gamma}/(1-\gamma)$ is concave and we only consider dual penalties $\pi(\mathbf{x})$ that are linear in $\mathbf{x}$, the inner problem is a convex optimization problem with linear constraints. Moreover, for the dual problem, we can significantly reduce the number of constraints in (2) by using as decision variables the units $\mathbf{n}_{t,t}$ purchased at time t, the units $\hat{\mathbf{n}}_{j,t} := \mathbf{n}_{j,t-1} - \mathbf{n}_{j,t}$ of cost-basis $\mathbf{p}_j$ sold at time t, and the cash account b_t. The new set of variables must satisfy the following constraints, which are equivalent to (2):

$$\begin{aligned} &b_t \geqslant 0, \quad \mathbf{n}_{t,t} \geqslant \mathbf{0}, \quad \hat{\mathbf{n}}_{j,t} \geqslant \mathbf{0}, \quad t \geqslant 0,\ 0 \leqslant j < t, \\ &\text{and} \quad \mathbf{n}_{t,t} = \sum_{j=t+1}^{T} \hat{\mathbf{n}}_{t,j} \quad t \geqslant 0. \end{aligned} \tag{32}$$

The inner maximization problem (31) then takes the following form[15]

$$\max_{b_t, \mathbf{n}_{t,t}, \hat{\mathbf{n}}_{j,t}, g_t^s, l_t^s, g_t^l, l_t^l, \hat{g}_t^l, \hat{l}_t^s} \left\{ \frac{b_T^{1-\gamma}}{1-\gamma} - \pi(\mathbf{x}) \right\}$$

s.t.

(9)–(15),
nonnegative constraints on trades (32),
boundary and initial conditions on the tax terms (5)–(6). (33)

Although (33) appears very similar to (8), there are significant differences. (8) is a stochastic optimization problem where the stock price process $\mathbf{p}_t$ is stochastic, and the decisions $(b_t, \mathbf{n}_{t,t}, \hat{\mathbf{n}}_{j,t})$ are stochastic processes adapted to the filtration $\mathcal{F}_t$. In contrast, (33) is a deterministic optimization problem corresponding to a single sample path $\mathbf{p}$. Note that there are $K(T+1)(T+2)/2+7T+4$ variables and $(K+3)(T+1)+3T$ constraints in this dual *inner* problem. Moreover, the constraints in the optimization problem (33) are quite sparse; only a small subset of the decision variables appears in each of the constraints, and the objective function in (33) is separable. Most convex optimization algorithms, including the MOSEK nonlinear convex optimization solver that we used, can take advantage of these two properties.

For our 25 stock and 80 time periods example, each dual problem instance has 83,589 decision variables and 2,508 constraints, and it can be solved in less than 0.5 seconds. Given that the dual problem size only grows quadratically in K and T, we can easily solve dual problem instances with a much larger number of securities and time periods. That these dual problems are so easy to solve is very surprising given the reputation of the (primal) tax problem for being an extremely challenging control problem to solve.

5. Numerical Experiments

We consider an exact cost-basis LUL problem where the agent can invest in 25 risky securities as well as the cash account. The investment horizon is 20 years with $T=80$ periods corresponding to a quarterly trading frequency. The U.S. tax code defines short-term gains (losses) to be those gains (losses) that are realized from selling securities that have been held for one year or less. All other gains and losses are long term. We therefore set $m_{\text{st}}=3$. We also assume short- and long-term capital gains tax rates of $\tau^s=40\%$ and $\tau^l=20\%$, respectively. The agent has power utility and we consider coefficients of relative risk aversion $\gamma\in\{1.5, 3, 5\}$.

5.1. Security Price Dynamics

We assume the net annual after tax risk-free interest rate is 1% so that we set $r_0=1.01^{1/4}$ per period. We assume the security return dynamics satisfy

$$\ln \mathbf{r}_t=\boldsymbol{\beta}_m\mathbf{f}_t^m+\boldsymbol{\beta}_{smb}\mathbf{f}_t^{smb}+\boldsymbol{\beta}_{hml}\mathbf{f}_t^{hml}+\boldsymbol{\epsilon}_t, \tag{34}$$

where $\mathbf{r}_t$ denotes the return vector over the period $[t-1,t]$, $\mathbf{f}_t^m$, $\mathbf{f}_t^{smb}$, and $\mathbf{f}_t^{hml}$ are the "market," "small minus big" and "high minus low" factors over $[t-1,t]$ defined by Fama-French three-factor model, and $\boldsymbol{\epsilon}_t$ are IID multivariate normal random vectors with mean $\mathbf{0}$ and covariance $\boldsymbol{\Sigma}_\epsilon$. We use the 25 portfolios formed on size and book to market as described on Ken French's website as the risky securities in our model. We obtain the factor loadings $\boldsymbol{\beta}_m$, $\boldsymbol{\beta}_{smb}$, and $\boldsymbol{\beta}_{hml}$ through a linear regression using data from 1988 to 2008 that is also available on his website. We also assume that the vector of factor returns $(\mathbf{f}_t^m, \mathbf{f}_t^{smb}, \mathbf{f}_t^{hml})$ is IID normally distributed across time periods. Therefore the return dynamics (34) do not[16] have a state vector $\mathbf{z}_t$, and so the solution to the no-tax problem of Appendix A.1 is such that it is optimal to pursue a myopic trading strategy that maximizes the expected utility of the next period's wealth.

5.2. Results

We simulated $I=5{,}000$ sample paths and implemented the three policies described in Section 3 along each sample path. The dual problem was also solved on the same 5,000 paths. Rather than reporting lower and upper bounds in terms of expected utilities which are difficult to interpret, we will instead report bounds as annualized certainty equivalent returns. This is a standard way to report results in dynamic portfolio optimization problems. Toward this end, let $\bar{u}$ denote the average utility of a feasible policy or the average of optimal dual objective functions across all sample paths. The certainty equivalent (CE) annualized return is then defined as the constant annualized return, r_{ce}, that yields the same average utility. In our power utility framework and with the setting that each period corresponds to three months this implies $(1/(1-\gamma))(w_0(1+r_{ce})^{T/4})^{1-\gamma}=\bar{u}$ which yields

$$r_{ce}=\left(\frac{((1-\gamma)\bar{u})^{1/(1-\gamma)}}{w_0}\right)^{4/T}-1. \tag{35}$$

We report the mean and 95% confidence intervals (CI) for r_{ce} in our numerical experiments. We used the realized utility of the no-tax model as a control variate to reduce the number of Monte Carlo paths that we required for estimating accurate primal and dual bounds. In particular, we estimate the expected utility $\bar{u}$ as

$$\bar{u}=\frac{1}{I}\sum_{i=1}^{I}\big(U\big(b_T(\mathbf{x}(\mathbf{p}^{(i)}))\big) + \beta\big[V_0^N(w_0)-U(\tilde{w}_T(\tilde{\mathbf{x}}(\mathbf{p}^{(i)})))\big]\big), \tag{36}$$

where I is the number of Monte Carlo paths, $\mathbf{p}^{(i)}$ is the sequence of security price vectors along the ith sample path, $\mathbf{x}(\mathbf{p}^{(i)})$ is the corresponding sequence of trade decisions made by the suboptimal policy or the optimal solutions of the dual problem under consideration, $\tilde{\mathbf{x}}(\mathbf{p}^{(i)})$ is the optimal feasible adapted trading policy for the no-tax problem and V_0^N is the optimal expected utility for the no-tax problem. We assumed an initial wealth of $w_0=1$, but, of course, the assumption of CRRA utility implies that our results do not actually depend on w_0. Numerical results are displayed in Table 1.

As expected, we see the tax blind policy performs very poorly and underperforms the other policies by more than 0.5%–2% per annum. In contrast, the tax-aware and RBH policies perform very well with the latter outperforming the former by only 1–6 basis points per annum. The upper bound provided by the no-tax solution is 20–60 basis points higher than the best lower bound, i.e., the CE return of the RBH strategy. The dual upper bound is much better, however. The gap between the RBH primal bound and the dual

Table 1. Annualized % certainty equivalent (CE) returns for the three feasible policies together with the dual and no-tax upper bounds.

	Lower bounds			Upper bounds	
Risk aversion (γ)	Tax blind	Tax aware	Rolling buy-and-hold	Dual	No-tax
1.5	3.25 (3.24, 3.27)	5.28 (5.28, 5.29)	5.34 (5.33, 5.34)	5.44 (5.44, 5.45)	5.95
3	2.20 (2.19, 2.22)	3.18 (3.17, 3.18)	3.22 (3.21, 3.22)	3.28 (3.28, 3.28)	3.54
5	1.74 (1.73, 1.74)	2.29 (2.28, 2.29)	2.30 (2.30, 2.31)	2.33 (2.33, 2.33)	2.51

Note. Corresponding approximate 95% confidence intervals are given in parentheses.

bound varies from 10 basis points when $\gamma = 1.5$ to just 3 basis points when $\gamma = 5$. This implies the RBH strategy is within at most 10 basis points of the optimal solution when $\gamma = 1.5$ and 3 basis points when $\gamma = 5$. For this data set at least, it seems reasonable to conclude that we have succeeded in approximately solving the tax-aware asset allocation problem. Moreover, given the problem features—LUL, exact cost-basis differential tax rates for short- and long-term gains, 80 time periods and 25 securities—we would argue that these numerical results are superior to other results in the literature.

5.3. Analysis of Risky Asset Holdings

It is of interest to consider the total investment made by the various trading strategies in the risky securities. Toward this end, we define the stock to wealth ratio, $(\sum_{j=0}^{t} \mathbf{p}_t' \mathbf{n}_{j,t})/(b_t + \sum_{j=0}^{t} \mathbf{p}_t' \mathbf{n}_{j,t})$, to be the percentage of pretax wealth invested in the risky securities. This ratio is plotted in Figure 1(a) over the entire investment horizon[17] for the various strategies and levels of risk aversion. The ratio for the no-tax optimal policy is horizontal as there are no state variables driving the security returns and so the no-tax optimal portfolio weights are constant across time. The tax-aware policy attempts to trade to the no-tax optimal portfolio so its ratio is very close to that of the no-tax optimal policy. In contrast, the RBH policy invests more than the no-tax optimal policy in the risky securities with the exception of the $\gamma = 1.5$ case when the ratio for all strategies is 100%. Not surprisingly, we see that the ratio decreases with risk aversion for all strategies. We can provide some intuition for why the RBG strategy invests more in the risky securities than the tax-aware policy. When the RBH strategy has unrealized gains, for example, then it can afford to take on more risk as some of the potential future losses can be used to offset the taxes that are due on the unrealized gains. This is clearly not true of the tax-aware heuristic policy because it simply attempts to trade to the no-tax solution, and therefore pays no attention to unrealized gains in determining its target portfolio.

In Figure 1(b), we plot the trading to wealth ratio, which we define to be $(\mathbf{p}_t' \mathbf{n}_{t,t} + \sum_{j=0}^{t-1} \mathbf{p}_t'(\mathbf{n}_{j,t-1} - \mathbf{n}_{j,t}))/(b_{t-1} r_0 + \sum_{j=0}^{t-1} \mathbf{p}_t' \mathbf{n}_{j,t-1})$ at time t. We note that the denominator is the pretrade portfolio wealth at time t while the numerator is the total dollar trading volume at time t. Several interesting observations can be made. First, for a given level of

Figure 1. (Color online) Stock to wealth and trading to wealth ratios as a function of t.

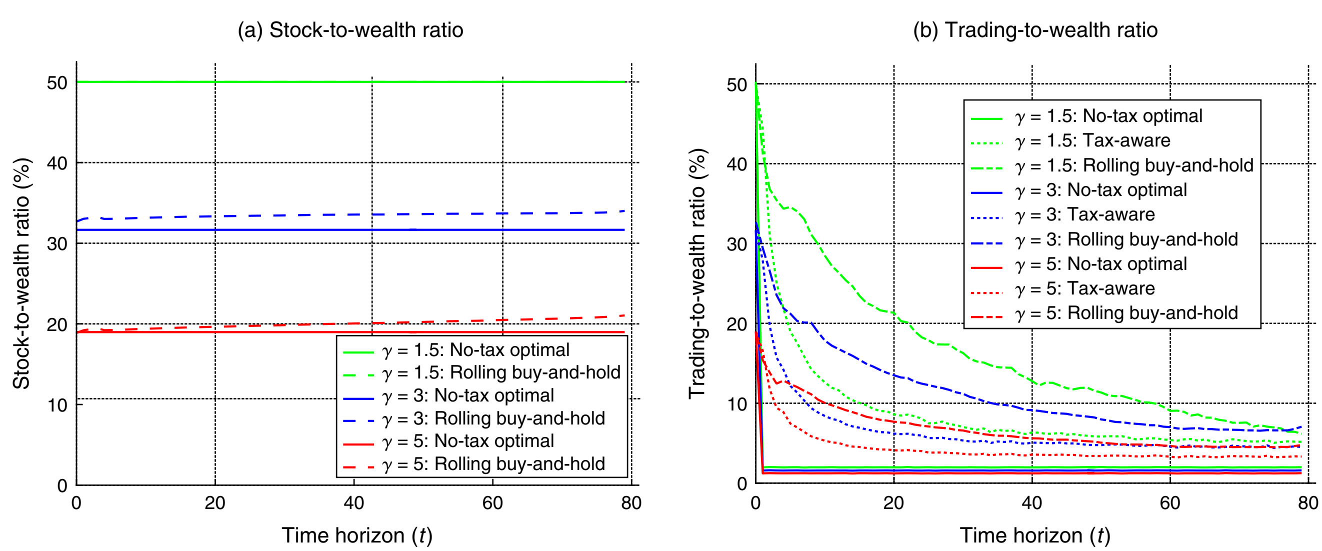

risk aversion we see that the RBH strategy trades a greater volume than the tax-aware policy, which, in turn, trades a greater volume than the no-tax optimal policy. This is not surprising since the RBH policy believes it can trade only once, i.e., in the current period, before having to liquidate the portfolio at maturity. The RBH policy is therefore more reactive to large price movements in any given period as it will then attempt to rebalance the portfolio (in a tax-efficient manner) in the belief that it has only one opportunity to do so. The tax-aware heuristic policy sees no such urgency and perceives little loss in failing to fully rebalance in any given period. It is also not surprising to see that the various policies trade more at lower levels of risk aversion. Finally, we also note that the trading to wealth ratio of the RBH and tax-aware policies gradually decrease in t.

The values in Figures 1(a) and 1(b) were obtained by averaging across the I simulated sample paths that were used to estimate the certainty equivalents of the policies. There is therefore some statistical noise in these numbers and this is clearly manifested in the figures.

5.4. Tax Forgiveness in the Final Period

The so-called step-up provision of the U.S. tax code allows an investor to bequeath assets upon his death and to have the cost basis of these assets automatically reset to the fair market value of the assets at the time of death. Given that most assets will generally have increased in value by the time of death, this is a considerable[18] tax advantage that essentially resets all capital gains tax rates to zero in the final period. In this subsection, we therefore rerun our numerical experiments but with $\tau^s = \tau^l = 0$ at time T. It is, of course, (typically) the case that the time of death is random but for now, it is convenient[19] to simply assume it is fixed and known to occur at time T. The results are displayed in Table 2.

Several observations are in order. First, and as expected, all three policies perform better with the step-up provision in place. The improvement in the tax blind policy is negligible, however, and ranges from only 1–3 basis points per annum. This is not surprising, however, since the tax blind policy is unaware of the step-up provision and therefore can only take advantage of it in the final period. In contrast, the RBH policy is aware at all times t of the step-up provision applying at time T. However, the more risk-averse investors will not be able to take full advantage of it, as doing so would require them to hold increasingly unbalanced portfolios as time T approaches. We therefore see the largest benefit accrue to the $\gamma = 1.5$ investor with increases of 58 basis points per annum for RBH policy.

The upper bounds are perhaps more interesting. The no-tax upper bound is, of course, unchanged from Table 1. The dual bound, however, has increased and appears to be identical to the no-tax upper bound. At least in this case then, the dual bound provides no information over and beyond that provided by the no-tax bound. However, this is not too surprising given that the gap from the best lower bound to the no-tax bound is already very small, ranging from just 3 basis points when $\gamma = 1.5$–5 basis points when $\gamma = 5$. It does raise the interesting question, however, as to whether it can be shown that dual and no-tax bounds are identical when the step-up provision is present. This is not, in fact, the case. While the same results of Brown and Smith (2011) (that we referred to in Section 4.1) imply that the dual bound is always at least as good as the no-tax solution, we observed in our numerical experiments that it yielded a small improvement over the no-tax bound on a small subset of dual problem instances. These differences were small and as they only occurred on a small subset of dual problem instances, the aggregate difference across all paths was very small and does not show up in the results of Table 2 where results are only reported to two decimal places.

We finally note in passing that the results of Brown and Haugh (2013) also provide some further insight on this observation. It is straightforward to show

$$V^N(x) \geqslant \mathcal{T}_{\mathrm{lul}} V^N(x), \tag{37}$$

$$V^N(x) \geqslant \mathcal{T}_{\mathrm{step}} V^N(x). \tag{38}$$

where V^N is the optimal value function for the no-tax problem, and $\mathcal{T}_{\mathrm{lul}}$ and $\mathcal{T}_{\mathrm{step}}$ are the Bellman operators for the corresponding LUL problem with and without the step-up provision in the final period, respectively. Moreover, standard results from dynamic programming show that any

Table 2. Annualized % certainty equivalent (CE) returns for the three feasible policies together with the dual and no-tax upper bounds.

	Lower bounds			Upper bounds	
Risk aversion (γ)	Tax blind	Tax aware	Rolling buy-and-hold	Dual	No-tax
1.5	3.28 (3.27, 3.30)	5.72 (5.71, 5.72)	5.92 (5.91, 5.93)	5.95 (5.95, 5.95)	5.95
3	2.22 (2.20, 2.23)	3.34 (3.34, 3.35)	3.49 (3.48, 3.50)	3.54 (3.54, 3.54)	3.54
5	1.75 (1.74, 1.75)	2.37 (2.36, 2.37)	2.46 (2.45, 2.47)	2.51 (2.51, 2.51)	2.51

Notes. Corresponding approximate 95% confidence intervals are given in parentheses. Step-up provision applies in the final period.

approximate value function that satisfies (37) and (38) must be an upper bound on the true optimal value functions, for the LUL problem with and without the step-up provision, respectively. Brown and Haugh (2013) show that the dual bounds constructed using an approximate value function that satisfies (37) and (38) are no worse than the bound provided by the approximate value function itself. In our context, this result simply states that our dual bounds must improve on the bound provided by the no-tax solution. Of course, we already know this using the aforementioned results from Brown and Smith (2011).

The significance of the Brown and Haugh (2013) treatment is that they provide a lower bound on the improvement of the dual bound over the bound provided by the approximate value function (which, in our case, is the no-tax solution V^N). In particular, their results suggest that the degree of bound improvement will be determined, in part, by the extent to which the inequalities in (37) and (38) are slack with less slack resulting in less bound improvement. While we expect there to be a lot of slack in (37) (and observe this to be the case in Table 1), there is considerably less slack in (38) and one can easily imagine states x, where there is essentially no slack. We are therefore not surprised[20] to see that the dual bound in Table 2 is more or less identical to the no-tax bound.

6. The Average Cost-Basis LUL Problem Formulation

Under the average cost-basis rule, the cost basis for any security in the portfolio is defined as the average price at which the shares of that security in the portfolio were purchased. The advantage of using the average cost basis is that at any time t, there is only one cost basis per security. In contrast, under the exact cost-basis approach, it is possible to have t separate tax bases for each security at each time t. Consequently, the average cost-basis formulation results in significantly fewer state variables and is therefore relatively easier to solve. The average cost-basis approach is also of considerable practical importance. For example, investors in U.S. mutual funds can often select an "average cost, single category" basis or "average cost, double category" basis instead of using the exact cost basis. The "average cost, double category" method requires the calculation of a separate average cost basis for long- and short-term holdings. This method is rarely used, however, because it is more complex to administer. The default approach for mutual funds is therefore the average cost, single category method where a single average cost basis is used for all holdings. In this section, we will consider the primal and dual versions of the average cost single category method. Moreover, we will assume a single tax rate that applies to all gains and losses.

Our goal here is twofold. First, we would like to investigate how much an investor loses by using the average cost basis instead of the more flexible exact cost-basis method. As suggested by DeMiguel and Uppal (2005) and others, the loss may well be very small and our goal is to confirm this in a much larger problem setting. Second, we shall see that dual problem instances for the average cost basis are nonconcave and are therefore much harder to solve than dual problem instances for the exact cost-basis problem. This result is somewhat surprising given that the primal average cost-basis problem, while still difficult, is much easier than the primal exact cost-basis problem. Because of the greater flexibility[21] provided by the exact cost-basis approach, we know the exact cost-basis dual bound is also an upper bound for the average cost-basis problem. But it is of independent interest to compare this upper bound with the average cost-basis dual bound, which will require the exact solution of nonconcave maximization problems. We will solve these latter problems using the polyhedral branch-and-cut approach of Tawarmalani and Sahinidis (2005). Showing that large nonconcave dual problem instances can be solved relatively quickly is a further contribution of this paper and should be of interest in other application domains where the dual problem instances are not easy to solve.

6.1. Problem Formulation

Let $s^+_{t,k}$ (resp. $s^-_{t,k}$) denote the number of shares of security k purchased (resp. sold) at time t. Let $s_{t,k}$ denote the number of shares in security k held after trading at time t, and let $\tilde{p}_{t,k}$ denote the associated average cost basis. Let $\mathbf{s}^+_t := [s^+_{t,1} \ \ldots \ s^+_{t,K}]'$, $\mathbf{s}^-_t := [s^-_{t,1} \ \ldots \ s^-_{t,K}]'$, $\mathbf{s}_t := [s_{t,1} \ \ldots \ s_{t,K}]'$, and $\tilde{\mathbf{p}}_t := [\tilde{p}_{t,1} \ \ldots \ \tilde{p}_{t,K}]'$. Using these variables, the LUL tax code can be simplified since all gains are subject to a single tax rate τ, and therefore there is no need to distinguish short- and long-term gains. In particular, we define

$$g_t = \max\{c_t + l_{t-1}, 0\} \quad \text{and} \quad l_t = \min\{c_t + l_{t-1}, 0\}, \tag{39}$$

where $c_t = (\mathbf{p}_t - \tilde{\mathbf{p}}_{t-1})'\mathbf{s}^-_t$ is the net realized proceeds from sales at time t. Note that (39) allows carried losses to offset gains and unused losses to be carried forward. We can now formulate the average cost-basis LUL problem:

$$\max_{(b_t, \mathbf{s}^+_t, \mathbf{s}^-_t, \mathbf{s}_t, g_t, l_t, \tilde{p}_{t,k}) \in \mathcal{F}_t,\, t=0,\ldots,T} \mathbb{E}_0\left[\frac{b_T^{1-\gamma}}{1-\gamma}\right] \tag{40}$$

s.t.

$$\mathbf{s}_t = \mathbf{s}_{t-1} + \mathbf{s}^+_t - \mathbf{s}^-_t, \quad \mathbf{s}_0 = \mathbf{s}^+_0,\ t \geqslant 1, \tag{41}$$

$$b_0 + \mathbf{p}'_0\mathbf{s}_0 = w_0, \tag{42}$$

$$b_t + \mathbf{p}'_t\mathbf{s}_t + \tau g_t = b_{t-1}r_0 + \mathbf{p}'_t\mathbf{s}_{t-1}, \quad t \geqslant 1, \tag{43}$$

$$g_t + l_t = (\mathbf{p}_t - \tilde{\mathbf{p}}_{t-1})'\mathbf{s}^-_t + l_{t-1}, \quad t \geqslant 1, \tag{44}$$

$$\tilde{p}_{t,k}s_{t,k} = \tilde{p}_{t-1,k}(s_{t-1,k} - s^-_{t,k}) + p_{t,k}s^+_{t,k}, \quad \forall k, t \geqslant 1, \tag{45}$$

$$b_t \geqslant 0, \quad \mathbf{s}^+_t \geqslant \mathbf{0}, \quad \mathbf{s}_{t-1} \geqslant \mathbf{s}^-_t \geqslant \mathbf{0}, \quad t \geqslant 0, \tag{46}$$

$$g_t \geqslant 0, \quad l_t \leqslant 0, \quad t \geqslant 1, \tag{47}$$

$$\tilde{\mathbf{p}}_0 = \mathbf{p}_0, \quad l_0 = 0, \tag{48}$$

and security price and market state dynamics.

Note that (41) updates the security positions after trading at time t, (42) (resp. (43)) is the budget constraint for $t=0$ (resp. $t>0$), (46) prohibits short sales or borrowing from the cash account, and (44) and (47) formulate the LUL tax rule (39). The average cost basis for each security is initially set to the time $t=0$ price and the initial carried losses are assumed to be 0 in (48). Suppose the agent buys $s^+_{t,k}>0$ shares and sells $s^-_{t,k}>0$ shares of security k at time t. Then, the net position is $s_{t,k}=s_{t-1,k}+s^+_{t,k}-s^-_{t,k}$, and the new cost basis is $\tilde{p}_{t,k}=(\tilde{p}_{t-1,k}(s_{t-1,k}-s^-_{t,k})+p_{t,k}s^+_{t,k})/s_{t,k}$. This rule is implemented in (45).

As before, we let $\mathbf{x}_t:=[b_t\ \mathbf{s}^{+\prime}_t\ \mathbf{s}^{-\prime}_t\ \mathbf{s}'_t\ g_t\ l_t\ \tilde{\mathbf{p}}'_t]$ denote the time-t decision variables. Similarly, we let

$$\mathbb{X}_0:=\big\{\mathbf{x}_0\mid \mathbf{x}_0\in\mathscr{F}_0 \text{ satisfies constraints (42), (46), and (48) at } t=0\big\},$$

$$\mathbb{X}_t(\mathbf{x}_{t-1},\tilde{\mathbf{p}}_{t-1}):=\big\{\mathbf{x}_t\mid \mathbf{x}_t\in\mathscr{F}_t \text{ satisfies constraints (41), (43)–(47)}\big\},\quad t\geqslant 1,$$

denote the sets of feasible trades at times $t=0$ and $t\geqslant 1$, respectively.

6.2. Suboptimal Policies

We use suitably adapted versions of the suboptimal policies of Section 3 for the average cost-basis problem. The Tax blind policy remains unchanged so that the agent liquidates her entire portfolio of risky securities at each time t, pays capital gains taxes on her net gains and then trades to the no-tax optimal portfolio weights. There is a slight change in the tax-aware policy because with just a single tax rate, there are no short-term capital gains taxes to avoid. More specifically, the agent still follows the five steps described in Section 3.2, but she simplifies steps (1), (3), and (4) by harvesting all losses, selling any shares as dictated by the no-tax solution and paying taxes as necessary, respectively.

It is easy to adapt the RBH strategy to the average cost-basis problem even though we now have product terms of decision variables in (44) and (45). The average cost RBH problem at time t takes the form:

$$\max_{\mathbf{x}_t,\,b_T^{(m)},\,g_T^{(m)}}\ \frac{1}{M}\sum_{m=1}^{M}\frac{(b_T^{(m)})^{1-\gamma}}{1-\gamma} \tag{49}$$

$$\begin{aligned}\text{s.t}\quad & \mathbf{x}_t\in\mathbb{X}_t(\mathbf{x}_{t-1},\tilde{\mathbf{p}}_{t-1}),\\ & b_T^{(m)}=b_t r_0^{T-t}+\mathbf{p}_T^{(m)\prime}\mathbf{s}_t-\tau g_T^{(m)},\quad 1\leqslant m\leqslant M,\\ & g_T^{(m)}=\max\{(\mathbf{p}_T^{(m)}-\tilde{\mathbf{p}}_t)'\mathbf{s}_t+l_t,0\},\quad 1\leqslant m\leqslant M,\end{aligned}\tag{50}$$

where M denotes the number of time T scenarios used to approximate the time T uncertainty. The average cost-basis $\tilde{\mathbf{p}}_{t-1}$ is known before trading at time t, and is, therefore not, a decision variable here. Hence, (44) that is included in the constraint $\mathbf{x}_t\in\mathbb{X}_t(\mathbf{x}_{t-1},\tilde{\mathbf{p}}_{t-1})$ is a linear constraint. The quadratic term $\tilde{\mathbf{p}}'_t\mathbf{s}_t$ in (50) can be linearized by replacing it by the right-hand side $\tilde{\mathbf{p}}'_{t-1}(\mathbf{s}_{t-1}-\mathbf{s}^-_t)+\mathbf{p}'_t\mathbf{s}^+_t$ of (45). Next, applying the smoothing technique introduced in Section 3.3, we approximate $g_T^{(m)}$ by

$$\tilde{g}_T^{(m)}(\mathbf{x}_t)=\frac{1}{\theta}\ln\big[\exp(\theta((\mathbf{p}_T^{(m)}-\tilde{\mathbf{p}}_{t-1})'(\mathbf{s}_{t-1}-\mathbf{s}^-_t)+(\mathbf{p}_T^{(m)}-\mathbf{p}_t)'\mathbf{s}^+_t+l_t))+1\big].\tag{51}$$

With these modifications, (49) can be approximated as follows:

$$\max_{\mathbf{x}_t\in\mathbb{X}_t(\mathbf{x}_{t-1},\tilde{\mathbf{p}}_{t-1})}\frac{1}{M(1-\gamma)}\cdot\sum_{m=1}^{M}\big(b_t r_0^{T-t}+\mathbf{p}_T^{(m)\prime}(\mathbf{s}_{t-1}+\mathbf{s}^+_t-\mathbf{s}^-_t)-\tau\tilde{g}_T^{(m)}(\mathbf{x}_t)\big)^{1-\gamma}.\tag{52}$$

This problem has a total of $2K+3$ variables: b_t, $\mathbf{s}^+_t$, $\mathbf{s}^-_t$, g_t, and l_t. Note that the number of decision variables in (52) does not increase with t, and is considerably fewer than that in the exact cost-basis problem (27). Furthermore, ignoring nonnegativity constraints, there are only two remaining constraints, namely, (43) and (44). In our numerical experiments, we found the average cost-basis RBH problem could be solved approximately three times as quickly as the corresponding exact cost-basis RBH problem.

6.3. Dual Bound

Following Proposition 1 in Appendix A.1, we know the value function of the no-tax problem is a valid upper bound for the average cost-basis LUL problem (40). However, this bound is likely to be too conservative as we saw in the exact cost-basis case of Section 5. While the exact cost-basis dual bound is a valid upper bound for the average cost-basis problem (and possibly a very good one), we would still like to consider the dual problem instances corresponding to the average cost-basis problem. These dual problem instances take the form

$$\max_{\mathbf{x}}\ \left\{\frac{b_T^{1-\gamma}}{1-\gamma}-\pi(\mathbf{x})\right\}\quad \text{s.t. constraints (41)–(48),}\tag{53}$$

where the sample path of prices $\mathbf{p}$ is known, and the average cost-basis $\tilde{p}_{t,k}$, trades $s^-_{t,k}$ and $s_{t,k}$, are all decision variables. Thus (44) and (45) are now nonconvex quadratic constraints, and the inner optimization problem that we need to solve to compute the average cost upper bound is no longer concave. Recall that we need[22] to compute the *global* optimum of this dual optimization problem to obtain a valid dual bound. Since it is hard, in general, to solve a nonconcave maximization problem, tackling (53) directly does not lead to an efficient strategy for computing provably valid dual bounds.

We have a valid dual bound for the average cost-basis LUL problem as long as we can *upper bound* (53) along each dual sample path. Suppose then that we simulate I dual sample paths and let $V_{\mathrm{up}}^{(i)}$ be the optimal solution of the dual problem (53) on the ith path. Let $I^o \subseteq I$ be those dual paths on which we can compute $V_{\text{up}}^{(i)}$, and let $I^u := I - I^o$ be the paths on which we can only obtain an upper bound, $\bar{V}_{\text{up}}^{(i)}$, for $V_{\text{up}}^{(i)}$. Since

$$\frac{1}{I}\sum_{i=1}^{I} V_{\text{up}}^{(i)} \leqslant \frac{1}{I}\left(\sum_{i\in I^o} V_{\text{up}}^{(i)} + \sum_{i\in I^u} \bar{V}_{\text{up}}^{(i)}\right), \tag{54}$$

the right-hand side of (54) is still a valid dual bound. Our approach, outlined below, produces such a bound. Moreover, we shall see that the quality of this bound remains very good because, at least in our numerical experiments, $\bar{V}_{\text{up}}^{(i)}$ is very close to $V_{\text{up}}^{(i)}$ for the paths in I^u.

6.3.1. Solving with BARON. Tawarmalani and Sahinidis (2005) proposed a polyhedral branch-and-cut approach to perform global optimization for nonconvex optimization problems. Their algorithm generated polyhedral cutting planes and relaxations for multivariate nonconvex problems, and their algorithm was implemented in the BARON solver (Sahinidis 2013).

Explicitly imposing tight lower and upper bounds for decision variables can significantly improve the efficiency of the BARON solver. We can take advantage of this as follows. First, recall that $\mathbf{r}_t$ denotes the time t return vector and that $\mathbf{r}_t$, for $t = 1, \ldots, T$, is known on each dual sample path. If the agent starts with an initial wealth of w_0, then it is clear that

$$\bar{w}_t = w_0 \prod_{j=1}^{t} \max(r_0, \mathbf{r}_j) \tag{55}$$

provides an upper bound on the time t wealth of the agent, where the max operator in (55) returns the maximum of r_0 and all elements in the return vector $\mathbf{r}_j$. We also note that the average cost basis for each security at time t must lie between the lowest and highest prices of that security along the sample path up to time t. Given these two observations, we can impose the following additional constraints for the decision variables along each dual sample path:

$$\mathbf{s}_t^+ \leqslant \bar{w}_t/\mathbf{p}_t, \quad \mathbf{s}_t^- \leqslant \bar{w}_t/\mathbf{p}_t, \quad \mathbf{s}_t \leqslant \bar{w}_t/\mathbf{p}_t, \quad g_t \leqslant \bar{w}_t,$$
$$-\bar{w}_{t-1} \leqslant l_t, \quad \min(\mathbf{p}_0, \ldots, \mathbf{p}_t) \leqslant \tilde{\mathbf{p}}_t \leqslant \max(\mathbf{p}_0, \ldots, \mathbf{p}_t). \tag{56}$$

We found that imposing these constraints resulted in a significant speedup of the BARON solver and allowed us to solve reasonably large average cost dual problem instances.

6.4. Numerical Experiments

In our numerical experiments, we used the security price dynamics as described in Section 5. We assumed a single tax rate of $\tau = 30\%$ and we assumed a total of 20 time periods. We restricted ourselves to 20 periods since solving the nonconcave dual problem instances with BARON was quite time consuming even after imposing the additional bounds on the variables. We implemented each of the three suboptimal policies along each simulated path and also solved the dual problem using the average cost basis and BARON. The results are shown in Table 3 together with results for the corresponding[23] exact cost-basis problem. We set the solving time limit for each dual path to be one minute and we used 2,000 paths in our numerical experiments. The global optimum on approximately 40% of these paths was found with a relative error of 10^{-6} within one minute. For the remaining 60% of the paths, we used the best upper bound that was returned by BARON when it stopped after reaching the one-minute cutoff.

There are two main observations to be made from Table 3. First, we note that the lower bound for the exact cost-basis and average cost-basis problems are almost identical. The only difference between the two (when CE returns are reported to two decimal places) is that the RBH

Table 3. Annualized % certainty equivalent (CE) returns for the three feasible policies together with the dual and no-tax upper bounds.

		Lower bounds			Upper bounds	
Risk aversion (γ)	Cost basis	Tax blind	Tax aware	Rolling buy-and-hold	Dual	No-tax
1.5	Exact	3.17 (3.15, 3.20)	4.27 (4.26, 4.28)	4.32 (4.28, 4.32)	4.61 (4.60, 4.63)	5.95
	Average	3.17 (3.15, 3.20)	4.27 (4.26, 4.28)	4.31 (4.28, 4.32)	4.62 (4.60, 4.63)	
3	Exact	2.07 (2.08, 2.08)	2.64 (2.63, 2.65)	2.70 (2.68, 2.72)	2.82 (2.81, 2.83)	3.54
	Average	2.07 (2.08, 2.08)	2.64 (2.63, 2.65)	2.69 (2.67, 2.71)	2.82 (2.80, 2.83)	
5	Exact	1.66 (1.65, 1.66)	1.97 (1.97, 1.98)	2.02 (2.68, 2.72)	2.08 (2.07, 2.08)	2.51
	Average	1.66 (1.65, 1.66)	1.97 (1.97, 1.98)	2.01 (2.67, 2.71)	2.08 (2.07, 2.08)	

Note. Corresponding approximate 95% confidence intervals are given in parentheses.

policy for the average cost-basis problem has a CE return that is one basis point lower than the RBH policy for the exact cost-basis problem. This is consistent with the observation of DeMiguel and Uppal (2005) who noted (in the one or two stock case) that the CE loss in wealth was small when using the average cost basis instead of the exact cost basis. Second, we note that the dual bounds[24] for the average cost-basis problems are also almost identical to the dual bounds of the corresponding exact cost-basis problems. The only noticeable difference is a one basis point difference when $\gamma = 1.5$.

We can draw three conclusions from these results. First, when solving tax-aware asset allocation problems with a single tax rate, there appears to be essentially no loss in restricting ourselves to the use of the average cost basis. Moreover, one can easily compute tight and valid upper bounds for the average cost-basis problem by using the corresponding exact cost-basis dual bounds, which are easy to compute. Second, the polyhedral branch-and-cut approach of Tawarmalani and Sahinidis (2005) can be employed successfully to obtain good dual bounds even when the dual inner problems are nonconcave, and are therefore difficult to solve. This is particularly the case when path-specific bounds on the variables are added to the problem formulation. We believe this approach may be useful in other application domains.

Finally, we note that the duality gaps in Table 3 are significantly wider than the corresponding gaps in Table 1. In particular, the gap[25] between the best primal policy, i.e., the RBH policy, and the dual bound ranges from approximately 29 basis points when $\gamma = 1.5$–6 basis points when $\gamma = 5$. The main explanation for this is that we used a single tax rate of $\tau = 30\%$ here so that all gains are taxed at this rate in Table 3. In contrast, the tax aware and (presumably) the RBH strategies always avoid short-term gains in the numerical results of Table 1, and therefore only paid taxes at the long-term rate of 20% there. This suggests that the no-tax solution is actually closer to the optimal solution (and therefore yields a better dual penalty) in Table 1 than in Table 3. This is not the complete story, however, as we can see by comparing the performance of the tax blind strategy in Tables 1 and 3. The tax blind strategy of Table 1 always pays taxes at a rate of 40% and yet still outperforms the tax blind strategy of Table 3 where the tax rate is smaller at 30%. This can be explained, in part, by the fact that the time horizon of 20 periods that we used here is considerably smaller than the 80 periods we used in Table 1. The shorter horizon allows less opportunity to use capital losses that have been carried forward. Moreover, the values of γ that we used here are not exactly comparable[26] to the corresponding values we used in Section 5 because of this difference in time horizons.

7. Conclusions and Further Research

In this paper, we have considered the challenging problem of tax-aware dynamic asset allocation. We have developed suboptimal and heuristic trading policies for the exact and average cost-basis LUL problems and constructed tight lower and upper bounds on the annualized certainty equivalent returns of these policies. It is clear that similar policies can also be constructed for other variations of these tax problems. Our principal contribution has been to demonstrate that much larger problems than previously considered can now be tackled through the use of sophisticated optimization techniques and duality methods based on information relaxations. Specifically, the dual formulations of exact cost-basis problems are much easier to solve than the corresponding primal problems and it is quite straightforward to solve dual problems with a realistic, i.e., large, number of securities and time periods. To the best of our knowledge, we are also the first to successfully use these duality methods for tax-aware asset allocation. We also consider the relatively easier average cost-basis problem in this paper but note that dual problem instances in this case are nonconvex. We propose solution approaches for these problem instances so that we can still obtain valid upper bounds. These solution approaches should be useful in other applications where the dual problem instances can be nonconvex.

There are several possible directions for future research. One direction is to improve our understanding of the (approximately) optimal trading policies for these problems. In particular, when do these policies trade and is there an easy to characterize no-trade zone? How do the answers to these questions depend on the particular version of the tax problem that we are addressing? A second direction is to apply the dual methodology to other tax-related problems such as tax-aware index tracking. One issue that then arises is the need to find good dual penalties. In this paper, we found that the no-tax optimal solution yielded a good dual penalty for the tax-aware asset allocation problem but an alternative penalty would need to be found for the tax-aware index tracking problem. It may also be the case that the no-tax solution does not yield sufficiently good dual penalties for some of the extensions outlined in the online appendix. In that case, we would again need to construct and evaluate better penalties. We are currently working on some of these problems.

Supplemental Material

Supplemental material to this paper is available at http://dx.doi.org/10.1287/opre.2016.1517.

Acknowledgments

The authors thank an anonymous associate editor and two anonymous referees for many helpful comments on an earlier version of the paper. The authors also thank David Brown and Jim Smith for their helpful comments.

Appendix A.

A.1. The No-Tax Problem Formulation

In the no-tax problem, capital gains are not taxed so that $\tau^s = 0$ and $\tau^l = 0$ for $t = 0, 1 \ldots, T$. The portfolio optimization problem can then be solved using dynamic programming (DP) as long as the dimension of the exogenous state vector $\mathbf{z}_t$ is sufficiently small.

Let $V_t^N(w_t, \mathbf{z}_t)$ denote the time-t optimal value function for the no-tax problem. Although the current wealth w_t is a state variable, it is well known that wealth can be factored out when the utility function belongs to the CRRA class as is the case in this paper. This means we can write

$$V_t^N(w_t, \mathbf{z}_t) = \frac{1}{1-\gamma} w_t^{1-\gamma} \phi_t(\mathbf{z}_t), \tag{57}$$

where $\phi_t(\mathbf{z}_t)$ is recursively defined with $\phi_T(\mathbf{z}_T) = 1$,

$$\begin{aligned} &\frac{1}{1-\gamma}\phi_t(\mathbf{z}_t) \\ &\quad = \max_{\tilde{b}_t, \tilde{\mathbf{n}}_t} \mathbb{E}_t\left[\frac{1}{1-\gamma}(\tilde{b}_t r_0 + \mathbf{r}'_{t+1}\tilde{\mathbf{n}}_t)^{1-\gamma}\phi_{t+1}(\mathbf{z}_{t+1})\right] && (58)\\ &\quad\quad \text{s.t.} \quad \tilde{b}_t + \mathbf{1}'\tilde{\mathbf{n}}_t = 1, && (59)\\ &\quad\quad\quad\quad\;\; \tilde{b}_t \geqslant 0,\ \tilde{\mathbf{n}}_t \geqslant \mathbf{0}, && (60)\end{aligned}$$

and $\tilde{b}_t$ and $\tilde{\mathbf{n}}_t$ denote the posttrade *fractions* of wealth w_t invested in the cash account and risky securities, respectively. The time t conditional expectation in (58) is taken over the next period's return vector $\mathbf{r}_{t+1}$ and the state vector $\mathbf{z}_{t+1}$. The budget constraint is given by (59), and (60) are the no borrowing and no short sales constraints, respectively. In this case, the dimension of the state space is equal to the dimension of the state vector $\mathbf{z}_t$, and so we can solve the no-tax optimal portfolio position $(\tilde{b}_t^*, \tilde{\mathbf{n}}_t^*)$, which are the optimal solution of (58) and V_t^N numerically over a fine grid of possible values of $\mathbf{z}_t$ if the dimension of $\mathbf{z}_t$ is sufficiently small. We can then derive heuristic trading policies for the tax problem based on $(\tilde{b}_t^*, \tilde{\mathbf{n}}_t^*)$, and also use V_t^N as a basis for constructing dual feasible penalties as described in Appendix A.3 below.

The following proposition compares the set of feasible policies for the exact or average cost-basis LUL problem to those for the no-tax problem. Note that the result in Proposition 1 does *not* hold, in general, for FUL problems.

Proposition 1. *The set of feasible policies for the exact or average cost-basis LUL problem is a subset of the set of feasible policies for the no-tax problem. Therefore* $V_t^N \geqslant V_t^{\text{LUL}}$.

Proof. Suppose taxes in the LUL problem are not paid to the tax authority; instead, they are invested in a special risk-free security that also earns the same risk-free rate r_0. The proceeds of this investment are returned to the agent at maturity T. In this case, we see that any feasible policy for this adjusted (exact or average cost-basis) LUL problem is also feasible for the no-tax problem. In particular, the optimal value function for this adjusted problem $V_t^{\text{LUL, adj}}$ satisfies $V_t^{\text{LUL, adj}} \leqslant V_t^N$. But we clearly have $V_t^{\text{LUL}} \leqslant V_t^{\text{LUL, adj}}$ and so the result follows. ☐

While an obvious result, the significance of Proposition 1 is that it guarantees dual bounds constructed using the no-tax optimal solution that will be at least as good as the no-tax bound itself. This follows from the results in Brown and Smith (2011). See Appendix A.3 for an additional discussion on this observation.

A.2. Further Details on the Tax-Aware Policy

We provide some brief examples here that illustrate the workings of the tax-aware policy for the case where there is a single risky security. Suppose after harvesting losses, the agent has $l^s = \$50$ of realized short-term losses and $l^l = \$100$ of realized long-term losses available to her. There is just a single risky stock and its current price is \$10. The agent has 100 long-term shares with cost-basis \$8 and 100 short-term shares with cost-basis \$9. We assume the long-term tax rate is 20%.

Case 1. Suppose now she wants to sell 40 shares of the stock (to achieve the optimal portfolio weights of the no-tax problem). She will sell 40 long-term shares and realize $(10-8)\times 40 = \$80$ of long-term gains. The agent will use \$80 of her realized long-term losses to offset these gains. The net long-term gains are therefore $\$80 - \$80 = \$0$, she does not pay any taxes in this period and she carries the remaining \$50 of short-term losses and \$20 of long-term losses forward to the next period.

Case 2. She wants to sell 80 shares. She will sell 80 long-term shares and realize $(10-8)\times 80 = \$160$ of long-term gains. The agent will first use \$100 of her realized long-term losses and then \$50 of her realized short-term losses to offset these gains. The net long-term gains are then $160-100-50 = \$10$, and the agent pays long-term capital gains taxes of $\$10 \times 0.2 = \$2$. No losses are carried forward.

Case 3. She wants to sell 120 shares. She will first sell all 100 long-term shares and realize $(10-8)\times 100 = \$200$ of long-term gains. She will use \$100 of her realized long-term losses as an offset after which the remaining long-term gains are $200 - 100 = \$100$. She then sells 20 short-term shares, which realize $(10-9)\times 20 = \$20$ of short-term gains which she offsets with \$20 of her realized short-term losses. She still has \$30 of excess short-term losses, however, so she will use them to offset some of her remaining long-term gains. This leaves her with net long-term gains of $100 - 30 = \$70$. She pays long-term capital gains taxes of $70 \times 0.2 = \$14$ and no losses are carried forward to the next period.

Case 4. She wants to sell 180 shares. She will first sell all 100 long-term shares and realize $(10-8)\times 100 = \$200$ of long-term gains. She will use \$100 of her realized long-term losses as an offset after which the remaining long-term gains are $200 - 100 = \$100$. She would like to sell an additional 80 shares but, in fact, she will only sell 50 short-term shares. This sale realize $(10-9)\times 50 = \$50$ of short-term gains, which are offset entirely by her \$50 of realized short-term losses. The agent then pays long-term capital gains taxes of $100 \times 0.2 = \$20$ and no losses are carried forward. She did not succeed in selling 180 shares because of the principle to avoid paying short-term capital gains taxes at all times.

A.3. Constructing Dual Feasible Penalties

For the exact and average cost-basis LUL problems, we consider a gradient penalty of the form originally proposed by Brown and Smith (2011). Recall that we use $\mathscr{X}^L$ to denote the set of feasible adapted policies for the exact cost-basis LUL problem. We also let $\mathscr{X}^N$ denote the set of feasible adapted policies for the no-tax problem. Given a sample path of security prices $\mathbf{p}$ and market states $\mathbf{z}$, let $\tilde{\mathbf{x}}^*(\mathbf{p}, \mathbf{z})$ denote the optimal feasible adapted trading policy for the no-tax problem, and let $\tilde{w}_T(\tilde{\mathbf{x}}^*(\mathbf{p}, \mathbf{z}))$ denote the corresponding terminal wealth that results from following this policy. We define our gradient penalty as follows:

$$\pi(\mathbf{x}) := \nabla_{\mathbf{x}} U(\tilde{w}_T(\tilde{\mathbf{x}}^*(\mathbf{p}, \mathbf{z})))'(\mathbf{x} - \tilde{\mathbf{x}}^*(\mathbf{p}, \mathbf{z})), \tag{61}$$

where $\nabla_{\mathbf{x}} U(\tilde{w}_T(\mathbf{x}))$ denotes the gradient of the terminal utility with respect to $\mathbf{x}$. Note that $\pi(\mathbf{x})$ is clearly linear in the decision vector $\mathbf{x}$. Since we can view the no-tax problem of (57) as a static convex optimization problem over the trading *policy* $\mathbf{x}$, the first-order conditions imply that

$$\mathbb{E}\left[\nabla_{\mathbf{x}} U(\tilde{w}_T(\tilde{\mathbf{x}}^*(\mathbf{p},\mathbf{z})))'(\mathbf{x}-\tilde{\mathbf{x}}^*(\mathbf{p},\mathbf{z}))\right] \leqslant 0 \qquad (62)$$

for any feasible adapted trading policy $\mathbf{x} \in \mathscr{X}^N$. By Proposition 1, it follows that $\mathscr{X}^L \subseteq \mathscr{X}^N$ so that (62) also holds for all $\mathbf{x} \in \mathscr{X}^L$. It now follows easily from a similar argument to the discussion preceding Proposition 4.2 in Brown and Smith (2011) that the gradient penalty (61) is dual feasible for the exact cost-basis LUL problem and that the resulting dual bound is guaranteed to be no worse than the bound provided by the no-tax solution.

It is also worth mentioning that more recently, Brown and Smith (2014) identified an alternative gradient penalty that is related to (61). This penalty explicitly includes Lagrange multiplier terms for constraints, e.g., no short sales constraints, that are binding in the relaxed[27] problem. They show that this related penalty is superior to (61) for concave dynamic programs in the sense that the resulting dual bound is guaranteed to be at least as good as the dual bounds obtained from using (61). If the constraints are never binding, however, then the two penalties are identical, and therefore result in the same dual bound. Given the agent's risk aversion, we found the no short sales and no borrowing constraints for the no-tax problem to be rarely if ever binding in our calibrated Fama–French model. For that reason, we continued to use the penalty in (61) rather than the new gradient penalty proposed by Brown and Smith (2014). It is clear from the numerical experiments in Section 5 that this penalty performed extremely well and we would expect no discernible improvement had we used the alternative gradient penalty.

While the above discussion refers to the exact cost-basis LUL problem, we note that Proposition 1 also applies to the average cost-basis LUL problem and that the penalty of (61) is also dual feasible for this problem.

A.4. Further Details on Solving the Rolling Buy-and-Hold Problems

Generating LDS Scenarios. At each time period t, we generate $M = 10{,}000$ scenarios of time T security prices. Assuming the dynamics specified in (34), the distribution of $\mathbf{p}_T$ conditional on the information available at time t is available in closed form. We can therefore generate time T scenarios directly, i.e., in a single step, rather than having to simulate scenarios at intermediate time points using (34).

We used low discrepancy sequences (LDS) (Glasserman 2004) instead of naive Monte Carlo since LDS fill the "unit cube" more uniformly and often result in a dramatic improvement in performance. In our numerical experiments, we used a K-dimensional Sobol sequence to generate samples of the time T risky security prices. We produced these LDS using `Matlab`'s LDS functionality. We rejected the first 1,000 points, retained every 101-st point thereafter, and also applied the so-called Matousek-Affine-Owen scrambling scheme.

The SQP Algorithm. Let $f_t(\mathbf{x}_t)$ denote the objective function in (27). We solve it using the following sequential quadratic programming (SQP) approach:

1. Choose a starting point, $\bar{\mathbf{x}}_t = [\bar{b}_t\ \bar{\mathbf{n}}'_{0,t}\ \ldots\ \bar{\mathbf{n}}'_{t,t}\ \bar{g}^s_t\ \bar{l}^s_t\ \bar{g}^l_t\ \bar{l}^l_t\ \bar{\hat{g}}^s_t\ \bar{\hat{l}}^s_t]$.

2. Approximate $f_t(\mathbf{x}_t)$ with a second-order Taylor expansion $\bar{f}_t(\mathbf{x}_t)$ about $\bar{\mathbf{x}}_t$, and solve

$$\begin{aligned}\max_{\mathbf{x}_t}\ \ &\bar{f}_t(\mathbf{x}_t) := f_t(\bar{\mathbf{x}}_t) + \nabla f_t(\bar{\mathbf{x}}_t)'(\mathbf{x}_t - \bar{\mathbf{x}}_t)\\ &\qquad + \tfrac{1}{2}(\mathbf{x}_t - \bar{\mathbf{x}}_t)'\nabla^2 f_t(\bar{\mathbf{x}}_t)(\mathbf{x}_t - \bar{\mathbf{x}}_t) \qquad (63)\\ \text{s.t.}\ \ &\mathbf{x}_t \in \mathbb{X}_t(\mathbf{x}_{t-1}, \mathbf{p}_{0:t}),\end{aligned}$$

where $\nabla f_t(\bar{\mathbf{x}}_t)$ and $\nabla^2 f_t(\bar{\mathbf{x}}_t)$ are the gradient vector and Hessian matrix of f_t, respectively, evaluated at $\bar{\mathbf{x}}_t$. Recall that we have analytical expressions for $\nabla f_t(\bar{\mathbf{x}}_t)$ and $\nabla^2 f_t(\bar{\mathbf{x}}_t)$. Let $\mathbf{x}^{\text{opt}}_t$ be the optimal solution to (63).

3. Evaluate the objective function f_t at $\mathbf{x}^{\text{opt}}_t$, and stop if we have converged to within a given error tolerance. Otherwise, set $\bar{\mathbf{x}}_t = \mathbf{x}^{\text{opt}}_t$ and return to Step 2.

We used an absolute error tolerance of 10^{-5} in our SQP algorithm. Depending on the level of risk aversion, γ, this corresponds to a relative error tolerance between 10^{-5} and 10^{-6}.

A.5. Exact Cost-Basis Problem Formulation for the Single Tax Rate Case

The exact cost-basis asset allocation problem for the case of just a single tax rate can be formulated as

$$\max_{(b_t, \mathbf{n}_{j,t}, g_t, l_t) \in \mathscr{F}_t,\, t=0,\ldots,T} \quad \mathbb{E}_0\left[\frac{b_T^{1-\gamma}}{1-\gamma}\right] \qquad (64)$$

s.t.

$$b_0 + \mathbf{p}'_0 \mathbf{n}_{0,0} = w_0, \qquad (65)$$

$$b_t + \sum_{j=0}^{t} \mathbf{p}'_t \mathbf{n}_{j,t} + \tau g_t = b_{t-1} r_0 + \sum_{j=0}^{t-1} \mathbf{p}'_t \mathbf{n}_{j,t-1}, \quad t \geqslant 1, \qquad (66)$$

$$\mathbf{n}_{t,t} \geqslant \mathbf{0}, \quad t \geqslant 0, \qquad (67)$$

$$\mathbf{n}_{j,t-1} \geqslant \mathbf{n}_{j,t} \geqslant \mathbf{0}, \quad t \geqslant 1, j < t, \qquad (68)$$

$$b_t \geqslant 0, \quad t \geqslant 0, \qquad (69)$$

$$g_t + l_t = \sum_{j=0}^{t-1} (\mathbf{p}_t - \mathbf{p}_j)'(\mathbf{n}_{j,t-1} - \mathbf{n}_{j,t}) + l_{t-1}, \quad t \geqslant 1, \qquad (70)$$

$$g_t \geqslant 0, \quad t \geqslant 1, \qquad (71)$$

$$l_t \leqslant 0, \ l_0 = 0, \quad t \geqslant 1, \qquad (72)$$

and security price and state variable dynamics, where g_t denotes taxable capital gains after trading at time t and l_t represents the accumulated unused losses, which will be carried forward. It is easy to see that (64)–(72) are a reduced form of the constraints in Section 2 that we obtain by adding short- and long-term gains together and applying a single tax rate τ.

Endnotes

1. This strategy is sometimes referred to as "shorting the box." The 1997 Tax Reform Act in the U.S. rules out "shorting the box" transactions.
2. We note that their algorithm required 90 hours with 100 CPUs working in parallel to solve the problem. They do not provide processor or software details.
3. In the current U.S. tax code, it is also possible to use excess long-term losses to offset short-term gains that remain after offsetting by short-term losses. We do not model that feature here but we do briefly discuss how it could be modeled in Appendix B.

4. We use the "hat" notation here to denote intermediate quantities in our calculations. For example, $\hat{l}_t^s$, will be further adjusted to obtain a final value of the short-term losses l_t^s carried forward after trading at time t.
5. As mentioned earlier, this is consistent with the current U.S. tax code. We also show in an online appendix how excess long-term losses can be used to offset short-term gains. Indeed, the argument we used above to linearize (4) does not hold when excess long-term losses can be used to offset excess short-term gains.
6. While typically ruled out in practice, it is common in the literature to allow wash sales. This is advantageous for a portfolio with embedded capital losses because the wash sales transaction realizes the capital losses so they are immediately available for offsetting any realized gains. In the online appendix, we discuss how wash sales constraints can be imposed.
7. If the agent is able to realize all her desired sales to reach the target dollar position in Step (3) and pay no-taxes by offsetting all gains in Step (4), then she can purchase any needed securities to arrive at the no-tax optimal portfolio weights $(\tilde{b}_t^*, \tilde{\mathbf{n}}_t^*)$ in Step (5).
8. Obviously, other similar heuristics are possible. For example, rather than selling shares in proportion to their long-term (or short-term) holdings, we could adopt a first in-first out heuristic or a heuristic, which sells shares with the highest cost basis first.
9. In our numerical experiments of Section 5, we consider a problem with 80 periods corresponding to quarterly investing for 20 years. Modeling the U.S. tax code, we therefore take $m_{\mathrm{st}} = 3$.
10. Further details are provided in the appendix.
11. When we count the number of constraints here and in the later part of the paper, we exclude the nonnegativity constraints on the variables since these constraints have little impact on the solution time.
12. When γ is rational, (22) can, in fact, be reformulated as a second-order cone program (Alizadeh and Goldfarb 2003).
13. The computations were done with `Matlab` using the convex nonlinear optimization in MOSEK on a Windows 7 computer with 3.40 GHz Intel i7-4770 4 cores CPU and 8 GB of RAM. We used this computer for all computations reported in this paper.
14. Calculations are done with MOSEK's QP solver called from `Matlab`.
15. In problem (33), we also need to do variable substitution $\hat{\mathbf{n}}_{j,t} = \mathbf{n}_{j,t-1} - \mathbf{n}_{j,t}$ for constraints (10)–(12). We do not write them explicitly here to save space.
16. In unreported numerical results, we also allowed two of the factors to follow AR(1) dynamics. The resulting predictability (in the form of a two-dimensional state vector $\mathbf{z}_t$) from fitting this model was very modest, however, and made almost no difference to our numerical results despite the significantly increased computational cost of solving the no-tax problem.
17. We omitted the final time period T as all risky security holdings must be liquidated at that point.
18. This provision is also somewhat controversial with regular calls to have it removed from the U.S. tax code. It is perhaps also worth mentioning that there is also a step-down provision, whereby assets that have decreased in value at the time of death also have their cost basis reset to the current market value. These provisions therefore provide incentives for holding appreciated assets until after death and selling depreciated assets before death.
19. We show in Appendix B that random time horizons can be easily handled.
20. It is worth mentioning that the results in Brown and Haugh (2013) apply to dual bounds constructed using approximate value functions rather than the gradient-based penalties we used and discussed in Appendix A.3. Their results were derived in the context of infinite-horizon problems but these problems do include the class of finite-horizon problems as a special case.
21. Indeed, we can view average cost-basis strategies as a strict subset of exact cost-basis strategies.
22. Unless we choose to use the exact cost-basis upper bound, which, as previously mentioned, is a valid upper bound for the average cost-basis problem.
23. The exact cost-basis problem for the case of a single tax rate is formulated in Appendix A.5.
24. The dual penalty $\pi(\mathbf{x})$ for the average cost-basis problem was constructed using the same gradient penalty that we used for the the exact cost-basis problem.
25. If we assume the true optimal solution lies at the midpoint of the duality gap, then we are only 15 and 3 basis points from optimality when $\gamma = 1.5$ and 5, respectively. We would therefore argue that the RBH strategy in Table 3 is still close to the optimal solution.
26. To make them exactly comparable, we should have adjusted them by the appropriate deterministic discount factor.
27. The relaxed problem in this paper is the no-tax problem, and the only constraints we imposed there are no short sales and no borrowing constraints.

References

Alizadeh F, Goldfarb D (2003) Second-order cone programming. *Math. Programming Ser. B* 95(1):3–51.
Brown DB, Haugh MB (2013) Information relaxations bounds for infinite horizon Markov decision processes. Technical report, Columbia University, New York.
Brown DB, Smith JE (2011) Dynamic portfolio optimization with transaction costs: Heuristics and dual bounds. *Management Sci.* 57(10):1752–1770.
Brown DB, Smith JE (2014) Information relaxations, duality, and convex stochastic dynamic programs. *Oper. Res.* 62(6):1394–1415.
Brown DB, Smith JE, Sun P (2010) Information relaxations and duality in stochastic dynamic programs. *Oper. Res.* 58(4):785–801.
Constantinides G (1984) Optimal stock trading with personal taxes: Implications for prices and the abnormal January returns. *J. Financial Econom.* 13(1):65–89.
Dammon RM, Spatt CS, Zhang HH (2001) Optimal consumption and investment with capital gains taxes. *Rev. Financial Stud.* 14(3):583–616.
Dammon RM, Spatt CS, Zhang HH (2002) Diversification and capital gains taxes with multiple risky assets. Technical report, Carnegie Mellon University, Pittsburgh, PA.
Dammon RM, Spatt CS, Zhang HH (2004) Optimal asset location and allocation with taxable and tax-deferred investing. *J. Finance* 59(3):999–1037.
DeMiguel V, Uppal R (2005) Portfolio investment with the exact tax basis via nonlinear programming. *Management Sci.* 51(2):277–290.
Dybvig PH, Koo HK (1996) Investment with taxes. Technical report, Washington University, St. Louis, MO.
Ehling P, Gallmeyer M, Srivastava S, Tompaidis S, Yang C (2010) Portfolio choice with capital gain taxation and the limited use of losses. Technical report, BI Norwegian Business School, Oslo, Norway.
Gallmeyer MF, Srivastava S (2010) Arbitrage and the tax code. Technical report, University of Virginia, Charlottesville, VA.
Gallmeyer MF, Kaniel R, Tompaidis S (2006) Tax management strategies with multiple risky assets. *J. Financial Econom.* 80(2):243–291.
Garlappi L, Naik V, Slive J (2001) Portfolio selection with multiple assets and capital gains taxes. Technical report, University of British Columbia, Vancouver, Canada.

Glasserman P (2004) *Monte Carlo Methods in Financial Engineering*, Vol. 53 (Springer, New York).

Hakansson NH (1970) Optimal investment and consumption strategies under risk for a class of utility functions. *Econometrica* 38(5): 587–607.

Merton RC (1969) Lifetime portfolio selection under uncertainty: The continuous-time case. *Rev. Econom. Statist.* 51(3): 247–257.

Rogers LCG (2007) Pathwise stochastic optimal control. *SIAM J. Control Optim.* 46(3):1116–1132.

Sahinidis NV (2013) *BARON* 12.1.0: *Global Optimization of Mixed-Integer Nonlinear Programs*. User's Manual, http://www.minlp.com/downloads/docs/baron%20manual.pdf.

Samuelson P (1969) Lifetime portfolio selection by dynamic stochastic programming. *Rev. Econom. Statist.* 51(3):239–246.

Tahar IB, Soner HM, Touzi N (2010) Merton problem with taxes: Characterization, computation, and approximation. *SIAM J. Financial Math.* 1(1):366–395.

Tawarmalani M, Sahinidis NV (2005) A polyhedral branch-and-cut approach to global optimization. *Math. Programming* 103(2): 225–249.

Martin Haugh is an associate professor of Practice in the IEOR Department at Columbia University. His research interests include financial engineering, dynamic optimization and data-driven and operations research.

Garud Iyengar is a professor in the IEOR Department at Columbia University. His research interests include convex, integer and robust optimization, portfolio selection and revenue management.

Chun Wang obtained his Ph.D. in operations research from Columbia University in 2014. His research interests are in financial engineering and sub-optimal control and he currently works at KCG, a securities trading firm.