

Portfolio Investment with the Exact Tax Basis via Nonlinear Programming

Victor DeMiguel

London Business School, 6 Sussex Place, Regent's Park, London NW1 4SA, United Kingdom, avmiguel@london.edu

Raman Uppal

Center for Economic Policy and Research, and London Business School, 6 Sussex Place, Regent's Park, London NW1 4SA, United Kingdom, ruppal@london.edu

Computing the optimal portfolio policy of an investor facing capital gains tax is a challenging problem: because the tax to be paid depends on the price at which the security was purchased (the tax basis), the optimal policy is path dependent and the size of the problem grows exponentially with the number of time periods. Dammon et al. (2001, 2002, 2004), Garlappi et al. (2001), and Gallmeyer et al. (2001) address this problem by *approximating* the exact tax basis by the weighted average purchase price. Our contribution is threefold. First, we show that the structure of the problem has several attractive features that can be exploited to determine the optimal portfolio policy using the *exact* tax basis via nonlinear programming. Second, we characterize the optimal portfolio policy in the presence of capital gains tax when using the *exact* tax basis. Third, we show that the certainty equivalent loss from using the average tax basis instead of the exact basis is very small: it is typically less than 1% for problems with up to 10 periods, and this result is robust to the choice of parameter values and to the presence of transaction costs, dividends, intermediate consumption, labor income, tax reset provision at death, and wash-sale constraints.

Key words: portfolio choice; capital gains tax; optimization; nonlinear programming

History: Accepted by David Hsieh, finance; received February 10, 2002. This paper was with the authors 3 days for 1 revision.

1. Introduction

Our objective in this paper is to study the optimal dynamic portfolio policy in the presence of capital gains tax when using the *exact* tax basis. This is an important problem given that most investors face taxes on their stock holdings. Moreover, the magnitude of the capital gains tax is quite large—it typically ranges from 20% to 40%—and so is much larger than transactions costs, which are usually less than 1%.

In spite of the importance of this problem, the optimal portfolio policies in the presence of capital gains tax using the exact tax basis have been studied only by Dybvig and Koo (1996), with the analysis being limited to four periods and a single stock. The main reason for this is that computing the optimal portfolio policy of an investor subject to capital gains taxes is a challenging task. The difficulty is that the tax to be paid depends not only on the selling price, but also on the price at which the securities were purchased, that is, the tax basis. As a consequence, the optimal policy is path dependent and the size of the problem grows exponentially with the number of time periods.

Constantinides (1983) shows that if short selling were costless and unconstrained, the optimal policy

would be to realize all losses immediately and to defer all gains.¹ The investor optimally defers all gains because, with costless short selling, she prefers to sell short those securities with an embedded capital gain instead of selling them outright (referred to as “shorting against the box”), avoiding any tax payment, and thus capturing the time value of taxes. Also, an investor should realize all losses by selling any securities whose price falls below their tax basis to get a tax rebate, and then rebalance her portfolio by buying the securities at the current price; this practice is known as a wash sale.² In essence, costless short selling allows one to separate the portfolio problem from the tax timing problem. However, in practice short selling is not costless and it may also be prohibited by the tax authorities, and thus, in this paper

¹ Constantinides (1983) also assumes symmetric taxation of long-term and short-term capital gains. In another paper, Constantinides (1984) shows that with asymmetric taxation an investor may optimally realize gains.

² Although Constantinides (1983) assumes that wash sales are allowed, they are disallowed by the US tax code. In particular, an investor is not allowed to obtain a tax rebate when selling a security with an embedded capital loss if the same security is bought within 30 days of its sale.

we focus on the portfolio problem for the case where short sales are prohibited.

A popular approach adopted in the literature for dealing with the complexity of the portfolio problem in the presence of capital gains tax has been to use the weighted average purchase price as the tax basis to find an approximate solution to the problem with short-sale constraints; see for example, Dammon et al. (2001, 2002, 2004), Garlappi et al. (2001), and Gallmeyer et al. (2001). This is equivalent to forcing the investor to sell the same proportion of the shares she holds for each different tax basis whenever she sells stock. The resulting policies are obviously sub-optimal because it is always better to sell those shares with the highest tax basis first. However, the advantage of using this approximation is that it makes the problem path independent, and thus, allows one to solve problems with a large number of dates using *dynamic programming*.

Dybvig and Koo (1996) formulate the problem with the exact tax basis as a nonlinear program and they propose two different algorithms to solve it, but their algorithms allow them to solve the problem for only four periods and a single stock.

... we find that the numerical algorithm works nicely for $T \leq 4$ periods with a single binomial risky asset with our current computing capacity, and it is clear that even orders of magnitude of increase in computing power will not allow us to handle much larger problems. (Dybvig and Koo 1996, p. 11)

Our first contribution is to show that the dynamic portfolio problem formulated by Dybvig and Koo (1996) has several attractive features that can be exploited to solve larger problems using a standard *nonlinear programming* algorithm such as SNOPT (Gill et al. 2002).³ We consider both problems with more periods (as many as 10) or more stocks in which the agent can invest (up to four stocks when there are four periods and up to two stocks for the case where there are seven periods). Our approach also allows one to consider other features of the real world not considered by Dybvig and Koo such as transaction costs, dividends, intermediate consumption, labor income, tax reset provision at death, and wash-sale constraints.

Our second contribution is to characterize the optimal portfolio policy in the presence of capital gains tax when using the *exact* tax basis. We find that in

the presence of capital gains taxes the investor holds a substantially larger proportion of stock. This is partly because taxes effectively reduce the after-tax volatility of the stock and partly because of the value of the tax-timing option. In addition, in the presence of taxes the investor no longer holds a constant stock-to-wealth ratio. As is well known, this is because in the presence of taxes there is a trade-off between the gains from optimal diversification and the taxes incurred from trading to reach the optimal portfolio position. However, when wash sales are constrained, the investor with a capital loss buys more *and* sells more than an investor who is not wash-sale constrained; hence, the investor's stock-to-wealth ratio oscillates much more than in the absence of this constraint. We also find that the investor chooses the optimal portfolio strategy so as to minimize taxes, and consequently, under the optimal portfolio policy the total amount of taxes paid is very small. In our numerical simulations, we find that the tax-to-wealth ratio is always less than 2%, and is on average less than 1% for all time periods.

Note that many of the characteristics discussed above also apply to the policy obtained when approximating the tax basis with the weighted average purchase price, as in Dammon et al. (2001, 2002, 2004), Garlappi et al. (2001), and Gallmeyer et al. (2001). But, to the best of our knowledge, our work is the first to analyze the portfolio policy with the exact tax basis and confirm the similarities between this policy and the approximate policies in the literature.

The ability to solve larger scale problems using the exact tax basis also enables us to compare the optimal portfolio policy in the presence of capital gains tax when using the *exact* tax basis to the suboptimal portfolio policy obtained from approximating the exact tax basis with the weighted average purchase price, as in Dammon et al. (2001, 2002, 2004), Garlappi et al. (2001), and Gallmeyer et al. (2001). Our third contribution is to show that the certainty equivalent loss from using the average tax basis instead of the exact basis is very small: it is typically less than 1% for problems with up to 10 periods. This result is robust to the choice of parameter values and to the presence of transaction costs, dividends, intermediate consumption, labor income, tax reset provision at death, and wash-sale constraints.

The small difference between the exact and approximate solutions can be explained by the fact that the investor following the optimal investment policy rarely holds shares bought at more than one date, and consequently, the weighted average purchase price is very close to the exact tax basis. The intuition for why the investor following the optimal investment policy rarely holds shares with more than one tax basis is that when the stock price goes up, a risk-averse investor rarely purchases any additional shares

³ For equilibrium models in the presence of taxes that are set in a static environment, see Schaefer (1982) and Dammon and Green (1987); Basak and Gallmeyer (2003) study the effects on asset prices of differences across individuals in tax rates on dividends in a dynamic model and Basak and Croitoru (2001) analyze the equilibrium implications of introducing redundant securities in such a setting.

of stock because of diversification reasons; thus, the tax basis of the shares held after an increase in the stock price is the same as in the previous time period. On the other hand, when the stock price goes down, the investor usually undertakes a wash sale and thus resets the tax basis of all those shares with an embedded capital loss to the current stock price. The only circumstance under which an investor may hold shares with two different tax bases is when the stock price goes down, and she holds shares with a tax basis lower than the current stock price. In the presence of transactions costs and wash-sale constraints, there is a higher likelihood that the investor will hold shares with more than one tax basis. Similarly, when cash dividends are added to the model they require a constant rebalancing of the investor's portfolio, and hence, lead to an increase in the proportion of stock with secondary tax basis. The inclusion of labor income to the model also has the same effect. However, our results indicate that even in these cases the certainty equivalent loss from using the weighted average purchase price rather than the true tax basis is quite small for a wide range of parameter values.

The rest of this paper is organized as follows. In §2, we describe the basic model of portfolio selection in the presence of taxes on capital gains and the optimization problem faced by the investor. In §3, we discuss properties of the optimization problem that can be exploited to efficiently compute the optimal portfolio policy. The optimal portfolio policies and their comparison to the approximate policies based on the average tax basis are presented in §4. We conclude in §5.

2. The Model and Optimization Problem

In the first part of this section, we describe the basic model proposed by Dybvig and Koo (1996) and in the second part we explain how the model can be extended to allow for other features such as transaction costs, dividends, intermediate consumption, labor income, tax reset provision at death, wash-sale constraints, and the presence of multiple risky assets. We adopt the same notation as Dybvig and Koo.

2.1. The Basic Model

Dybvig and Koo (1996) consider an investor facing two investment opportunities: a risk-free asset and a risky asset (stock). The risk-free asset yields an after tax risk-free rate r per time period. The stock price at time t is P_t and its evolution is modeled as a binomial process. The investor is assumed to have an initial cash endowment of C_0 .

The investor has to choose C_t , the amount of cash held after time $t = \{1, \dots, T\}$, and $N_{s,t}$, the number

of shares bought at time $s = \{0, \dots, T\}$ and kept after trading at time $t = \{s, \dots, T\}$. The variables $N_{s,t}$ allow one to keep track of the tax basis of each share in the portfolio. They must satisfy the following constraints:

$$N_{t,t} \geq N_{t,t+1} \geq \dots \geq N_{t,T} \geq 0 \quad \text{for all } t, \quad (1)$$

which imply that shares are being sold over time and any new purchases are indexed by the later time of purchase. In addition, the last constraint ($N_{t,T} \geq 0$) rules out short selling.

The investor's preferences are given by a standard power utility function with relative risk aversion γ , and her objective is to maximize the expected utility of cash at time T :

$$\max_{C_t, N_{s,t}} E \left[\frac{C_T^{1-\gamma}}{1-\gamma} \right]. \quad (2)$$

The investor's policy must satisfy the budget constraints

$$C_t = C_{t-1}r - N_{t,t}P_t + \sum_{s=0}^{t-1} (N_{s,t-1} - N_{s,t})(P_t - \tau(P_t - P_s)) \quad \text{for all } t. \quad (3)$$

The first term on the right-hand side, $C_{t-1}r$, is the riskless after-tax return on cash. The next term, $N_{t,t}P_t$, is the cost of the shares bought at date t . The last term is the after-tax proceeds from shares sold, where τ is the tax rate, P_t is the current stock price, and P_s is the price at which the shares were purchased, that is, the tax basis.

2.2. Extensions to the Basic Model

We now extend the Dybvig and Koo (1996) model described above to allow for transaction costs, dividends, intermediate consumption, labor income, tax reset provision at death, wash-sale constraints, and the presence of multiple stocks. These extensions will allow us to explore how the optimal portfolio policies depend on these real-world features.

2.2.1. Transaction Costs. Proportional transaction costs may be considered by using the following budget constraint:

$$C_t = C_{t-1}r - N_{t,t}(1 + \kappa)P_t + \sum_{s=0}^{t-1} (N_{s,t-1} - N_{s,t})((1 - \kappa)P_t - \tau(P_t - P_s)),$$

where κ is the proportion of the price paid as transaction cost. Compared to the budget constraint in Equation (3), we see that the effect of transactions costs is to increase the cost of shares purchased and reduce the proceeds from shares sold.

2.2.2. Dividends. We model dividends at date t as being a constant proportion d of the share price at date t . In addition, we tax dividends at the same rate as capital gains. This leads to the following budget constraint:

$$C_t = C_{t-1}r + (1 - \tau) d P_t \sum_{s=0}^{t-1} N_{s,t-1} - N_{t,t} P_t + \sum_{s=0}^{t-1} (N_{s,t-1} - N_{s,t})(P_t - \tau(P_t - P_s)) \quad \text{for all } t.$$

2.2.3. Intermediate Consumption. In the basic model we described, it was assumed that investors wish to consume only at the terminal date T . However, investors may wish to consume at each date $t \leq T$. The basic model can be extended to include intermediate consumption by introducing the additional decision variables con_t (the investor's consumption at time $t = \{0, \dots, T\}$). Then, the objective is to maximize the expected utility of consumption:

$$\max_{C_t, \text{con}_t, N_{s,t}} E \sum_{t=0}^T \left[\beta^t \frac{\text{con}_t^{1-\gamma}}{1-\gamma} \right],$$

where β is the subjective discount rate. In addition, we modify the budget constraint in (3) to account for the effect of withdrawing funds to finance consumption:

$$C_t = C_{t-1}r - N_{t,t}P_t + \sum_{s=0}^{t-1} (N_{s,t-1} - N_{s,t})(P_t - \tau(P_t - P_s)) - \text{con}_t.$$

2.2.4. Labor Income. In addition to financial income received from the portfolio, an investor may also be receiving labor income. We model labor income as a proportion ℓ of the investor's pre-tax wealth. This leads to the following budget constraint:

$$C_t = C_{t-1}r + \ell \left(C_{t-1}r + \sum_{s=0}^{t-1} N_{s,t-1} P_t \right) - N_{t,t} P_t + \sum_{s=0}^{t-1} (N_{s,t-1} - N_{s,t})(P_t - \tau(P_t - P_s)).$$

2.2.5. Tax Forgiveness. To allow for the possibility that the tax code forgives any capital gains tax upon the death of the investor, we use the following budget constraint for the last date, which is a modified version of Equation (3):

$$C_T = C_{T-1}r - N_{T,T}P_T + \sum_{s=0}^{T-1} (N_{s,T-1} - N_{s,T})P_T.$$

2.2.6. Wash-Sale Constraints. The basic model assumes that wash sales are allowed; that is, the investor is allowed to sell all the shares whose tax basis is above the current stock price to get a tax rebate, and then

buy back the optimal number of shares at the current stock price. If the tax code prohibits wash sales, the model can be modified by introducing the following additional wash-sale constraints:

$$N_{t,t} \sum_{s=0}^{t-1} (N_{s,t-1} - N_{s,t}) = 0 \quad \text{for all } t.$$

These constraints imply that one cannot buy and sell shares at the same date t . Note that the wash-sale constraint is meaningful only for a model where the decision frequency is monthly. However, it makes more sense to interpret our model on an annual rather than monthly frequency. Nevertheless, we consider this constraint here to understand its effect.

2.2.7. More Than One Stock. So far, the model considers only the case where the investor can hold only one risky asset. To extend the model to the case where the investor can hold multiple risky assets, let $N_{j,s,t}$ be the number of shares of stock $j = \{0, \dots, J\}$ bought at time $s = \{0, \dots, T\}$ and kept at time $t = \{s, \dots, T\}$. The investor's objective function may be written as

$$\max_{C_t, N_{j,s,t}} E \left[\frac{C_T^{1-\gamma}}{1-\gamma} \right].$$

Likewise, the budget constraint for all t may be written as

$$C_t = C_{t-1}r - \sum_{j=1}^J \left[N_{j,t,t} P_{j,t} + \sum_{s=0}^{t-1} (N_{j,s,t-1} - N_{j,s,t}) \cdot (P_{j,t} - \tau(P_{j,t} - P_{j,s})) \right],$$

where $P_{j,t}$ is the price of the j th stock at time t . Finally, the short-sale constraints are

$$N_{j,t,t} \geq N_{j,t,t+1} \geq \dots \geq N_{j,t,T} \geq 0 \quad \text{for all } j, t.$$

3. Characteristics of the Optimization Problem

In this section, we identify several attractive properties of the dynamic portfolio problem described above that can be exploited to efficiently compute the optimal portfolio policy using a standard *nonlinear programming* algorithm such as SNOPT (Gill et al. 2002). Recognizing these properties allows one to solve problems with more periods (as many as 10) or more stocks in which the agent can invest (up to four stocks when there are four periods and up to two stocks for the case where there are seven periods) than the ones considered in Dybvig and Koo (1996). This also allows one to consider other features of the real world not considered by Dybvig and Koo such as transaction costs, dividends, intermediate consumption, labor income, tax reset provision at death, and wash-sale constraints.

The basic model proposed in Equations (1), (2), and (3) is a nonlinear optimization problem subject to linear constraints.⁴ Moreover, its objective is to maximize a concave function. Much is known about how to solve these problems. First, for concave problems, we know that there exist nonlinear programming algorithms that will converge to a global maximizer from any starting point. In addition, the linearity of the constraints implies that the iterates generated by the algorithm may always remain feasible.

Another advantageous feature of problem (1)–(3) is that its objective and constraint functions are sparse; that is, only a few of the decision variables appear in each of the constraint and objective functions. Hence, sparse numerical linear algebra techniques can be used to reduce the amount of storage and the algebraic computations needed to solve this problem.

Finally, note that the objective function (2) is separable; that is, the objective function is a summation of terms and only one decision variable appears in each term. Again, most optimization algorithms can take advantage of this feature.

The binomial tree employed to model the stock price, however, implies that the number of variables and constraints in problem (1)–(3) will increase exponentially with the number of time periods. In particular, one can show that for S stocks and T periods, the number of variables, equality constraints, and inequality constraints is

$$\begin{aligned} \text{variables} &= (T + 1)(S + 1)^{T+1}, \\ \text{equality constraints} &= \frac{(S+1)^{T+1} - 1}{S}, \\ \text{inequality constraints} &= \frac{(S(T+1)-1)(S+1)^{T+1} + 1}{S}. \end{aligned}$$

We can use the above expressions to get an idea of the size of this model. For instance, when there is a single stock and seven periods, there are 2,048 variables, 255 equalities, and 1,793 inequalities. For one stock and 10 periods, there are 22,528 variables, 2,047 equalities and 20,481 inequalities. Finally, for two stocks and seven periods, there are 52,488 variables, 3,280 equalities, and 49,208 inequalities.

Despite the large size of problem (1)–(3), its favorable features—concavity and separability of the objective, linearity of constraints, and sparsity—allow nonlinear programming algorithms to solve it efficiently (see Luenberger 1984, Nocedal and Wright 1999 for an introduction to nonlinear programming).

⁴ The basic model remains a nonlinear optimization problem subject to linear constraints for all of the suggested extensions except when wash-sale constraints are added. The wash-sale constraints are nonlinear, nonconvex constraints known in the nonlinear programming literature as complementarity constraints; see Luo et al. (1996).

In particular, using the nonlinear programming software SNOPT of Gill et al. (2002), we solve problems with up to 10 periods and one stock or seven periods and two stocks.⁵

Dybvig and Koo (1996), on the other hand, proposed two different algorithms to solve problem (1)–(3) but could only solve problems of up to four periods and one stock. Both of their algorithms transform the problem into an *unconstrained* optimization problem and then solve it.

Their first algorithm eliminates the budget and short-sale constraints by introducing a huge number of decision variables corresponding to all possible activities the investor may undertake. One such activity, for instance, would be to buy one share of stock at time $T - 2$, sell it at time $T - 1$ only if the stock price increases, and otherwise sell it at time T . There are two difficulties associated with this reformulation. One, the resulting objective function is no longer separable and sparse, and two, the number of activities the investor may undertake is huge, for instance, a problem with seven periods has 4.4×10^{22} variables. Even with today's computing capacity, it would be impossible to solve a problem with such a large number of variables. In contrast, the number of decision variables in problem (1)–(3) is only 2,048.

Their second algorithm transforms the original problem (1)–(3) into an unconstrained problem with a smaller number of decision variables. In particular, the only decision variables left after transforming the problem are the number of shares of stock held by the investor at each date and state. Unfortunately, Dybvig and Koo (1996) show that the resulting optimization problem is nonsmooth. Solving nonsmooth problems is much harder than solving smooth problems.⁶ Not surprisingly, they report numerical difficulties when trying to use this algorithm.

Thus, the transformed unconstrained problems proposed by Dybvig and Koo (1996) are more difficult to solve than the original linearly constrained problem. We believe it is easier and more efficient to apply a standard nonlinear programming algorithm, such as SNOPT (Gill et al. 2002), to solve problem (1)–(3) directly, as illustrated in the next section.

4. Properties of the Optimal Portfolio Policies

In this section, we use SNOPT to solve the problem of choosing the optimal portfolio policy in the presence of a tax on capital gains. We solve investment

⁵ Gupta and Murray (2000) use SNOPT to compute optimal portfolio policies for investors with other utility functions. Portfolio optimization problems can also be solved by using stochastic programming (see Birge and Louveaux 1997, DeMiguel and Nogales 2002).

⁶ See Clarke (1983) for an introduction to nonsmooth optimization.

problems with seven periods and one stock, seven periods and two stocks, and 10 periods and one stock.⁷

We use these solutions to analyze the characteristics of the optimal portfolio policies and examine the sensitivity of the portfolio policies to the investor's risk aversion, the mean and volatility of stock returns, the presence of transaction costs, dividends, intermediate consumption, labor income, tax reset provision at death, asymmetric taxation, and wash-sale constraints, and the number of stocks and time periods.

We also compare the optimal investment policy to the suboptimal policy obtained by using the weighted average purchase price as the tax basis. Herein, we call the optimal policy *true tax basis policy* and the suboptimal policy *average tax basis policy*.⁸ To better understand the magnitude of the difference between the true and average tax basis policies, we also compute the following suboptimal policies: (i) the buy-and-hold policy, (ii) the realize-all-capital-gains-and-losses policy, and (iii) the augmented-buy-and-hold policy, as suggested by Dammon et al. (2001). In the buy-and-hold policy, the investor can buy stock on the first date, but is not allowed to buy or sell stock ever after, except on the final date. In the realize-all-gains-and-losses policy, the investor can choose how much stock to hold at the end of each date, but she is forced to realize all capital gains and losses at the beginning of each date. Finally, the augmented-buy-and-hold policy is the buy-and-hold policy augmented to allow the investor to realize her losses to obtain tax rebates but hold on to the gains.

4.1. Optimal Portfolio Policy in the Basic Model with Benchmark Parameters

In this section, we first report the benchmark parameter values used for analyzing the basic model. We label the problem with the basic model and the benchmark parameter values the "base case." We then describe the optimal portfolio policy using the true tax basis for the benchmark parameter values, and conclude by comparing the optimal policy to the approximate policy obtained when using the average tax basis. Then, we examine the sensitivity of our results to the choice of parameter values in §4.2 and to the modeling assumptions in §4.3.

⁷ We also solve the problem with four periods and four stocks but the results for this case are not reported.

⁸ The basic model in Equations (1), (2), and (3) can be modified to compute the average tax basis policy by introducing the additional constraint $N_{s,t} = \phi_t N_{s,t-1}$ for all t and $s < t$. This is equivalent to forcing the investor to sell at each date t the same proportion ϕ_t of the shares she holds for each different tax basis. Note that these additional constraints are nonlinear constraints because the components of ϕ_t are decision variables. Thus, unlike with dynamic programming, the average tax basis policy is harder to compute with nonlinear programming than the true tax basis policy.

4.1.1. Benchmark Parameter Values. For the base case, we use an investment problem with a horizon of seven periods and one stock. The stock price is assumed to follow the multiplicative binomial process described in He (1990), with expected annual rate of return of 10% and annual volatility of 20%. The pre-tax risk-free rate is assumed to be 6%. We further assume for our base case that there are no transaction costs and the stock pays no dividends. We consider an investor who has power utility with a risk-aversion parameter $\gamma = 3$, whose initial endowment is \$1, does not receive any additional labor income, and only consumes at the final date. Finally, for the base case we also assume that there is no tax reset provision at death, that short and long-term capital gains are taxed at the same rate of 35%,⁹ and that wash sales are allowed.

4.1.2. The True Tax Basis Policy. We compute the true tax basis policy using SNOPT to solve the base case problem.¹⁰ The top panel in Figure 1 shows the stock-to-wealth ratio for the true tax basis policy for *each* of the 128 paths in the binomial tree for the first seven dates.¹¹ At $t = 1$ the stock-to-wealth ratio is 53% for all 128 paths. At $t = 2$, for those paths where the stock price goes up (paths 65–128) the ratio increases to 58%, while for those paths for which the stock price goes down (1–64) the stock-to-wealth stays at 53%. Each time the stock price increases, the proportion of stock-to-wealth increases because the tax on capital gains deters the investor from rebalancing optimally. On the other hand, when the stock price goes down, the stock-to-wealth ratio remains relatively unchanged; this is because the investor undertakes a wash sale that reduces substantially the effect of the tax on capital gains.

From the top panel of Figure 1, we see that the investor's stock-to-wealth ratio varies between 53% and 66%, whereas in the absence of taxes, an investor faced with the base case problem would hold a constant stock-to-wealth ratio of 36%.¹² Note that, as reported in Dammon et al. (2001), in the presence of capital gains taxes the investor holds a substantially larger proportion of stock. This is partly because


⁹ The long-term capital gains tax in the U.S. tax code is 20% but here we set it to be 35% to be consistent with other papers in this literature, for instance, Dammon et al. (2001); the case where short- and long-term capital gains are taxed at different rates is analyzed in §4.3.6.

¹⁰ SNOPT is accurate to eight significant digits of the objective function. This implies that the accuracy of the decision variables and the certainty equivalent is approximately four significant digits.

¹¹ The investor optimally sells all her stock at the last date to finance consumption.

¹² The optimal portfolio policy in the absence of taxes can be obtained by solving problem (1)–(3) with $\tau = 0\%$.

Figure 1 True Tax Basis Policy for the Base-Case Problem



Note. This figure graphs the true tax basis policy for the base-case problem. The top panel gives the stock-to-wealth ratio for each of the 128 paths in the binomial tree and for each of the first seven dates. The bottom panel gives the proportion of secondary stock; that is, the proportion of stock held at a tax basis different from the basis of the majority of the stock.

taxes effectively reduce the after-tax volatility of the stock and partly because of the value of the tax-timing option. In addition, note that in the presence of taxes the investor no longer holds a constant stock-to-wealth ratio. As is well known, this is because in the presence of taxes there is a trade-off between the gains from optimal diversification and the taxes incurred from trading to reach the optimal portfolio position.¹³

4.1.3. Comparison with Suboptimal Portfolio Policies. There are several dimensions along which one can compare the true tax basis policy with the four suboptimal policies: the average-basis policy (avgbase), the buy-and-hold policy (buyhold), the augmented buy-and-hold policy (augbuy), and the realize-all-capital-gains-and-losses policy (realize).

One way is to compute the certainty equivalent value resulting from a particular portfolio policy and then computing the percentage difference between the certainty equivalent under the optimal policy and the suboptimal policies.¹⁴ A second comparison is in terms of the proportion of shares purchased at different dates and the taxes paid under each portfolio strategy. A third way is to compare the actual portfolio policies. We compare below the optimal and suboptimal portfolio policies using these three measures.

The buy-and-hold, the realize-all-capital-gains-and-losses, and the augmented-buy-and-hold policies yield a certainty equivalent loss of 0.48%, 1.09%, and 0.08%, respectively, with respect to the true tax basis policy. Note that, while the certainty equivalent losses incurred by the buy-and-hold and the realize-all-capital-gains-and-losses are considerable, the loss incurred by the augmented-buy-and-hold policy is very small, similar to the finding in Dammon et al. (2001).

But what is more interesting is that the difference between the true and average policies in terms of certainty equivalent is negligible (0.00%). We find this result surprising. The difference between the true and average policies stems from the fact that when selling stock, the investor using the true tax basis can sell those shares with the highest tax basis first. The investor using the average tax basis, on the other hand, effectively sells the same proportion of the shares she holds for each different tax basis. One would expect that by selling the shares with highest tax basis first, considerable tax savings could be achieved. Our numerical results, however, show the opposite.

But an examination of the portfolio policies under the true tax basis and the policies under the average basis shows that there are two reasons why their certainty equivalents are almost identical. The first reason is that an investor following the optimal investment policy almost never holds shares bought at more than one date. This can be observed from the lower panel of Figure 1, which shows the proportion of stock held at a *secondary basis*; that is, the proportion of stock held at a basis different to the tax basis of most of the stock. Note that this proportion is below 10% at all dates and scenarios, and is lower than 1% on average. Consequently, there is little advantage in being able to sell those shares with the highest tax basis first, because most of the time all shares have the same tax basis. The explanation for this low proportion of secondary basis stock is that, when the stock price goes up, a risk-averse investor would rarely purchase any additional shares of stock

¹³ See Dammon et al. (2001) and Garlappi et al. (2001) for a detailed analysis of this trade-off.

¹⁴ We define the certainty equivalent as the inverse utility function of the expected utility; that is, certainty equivalent = $u^{-1}(Eu(w))$.

because of diversification reasons. Consequently, the tax basis of the shares held after an increase in the stock price is the same as in the previous time period. On the other hand, when the stock price goes down, to get a tax rebate the investor sells all those shares whose tax basis is above the current stock price, and then buys shares at the current stock price (this can be seen from the lower panel of Figure 1, where we can see for the first 64 paths the investor typically holds shares with only a single tax basis, and even when shares with a second tax basis are held the proportion is less than 4%). The only circumstance under which an investor holds shares with two different tax basis is when the stock price goes down, and she holds shares with a tax basis lower than the current stock price. In this case, the investor would usually keep all shares whose tax basis is lower than the current stock price to avoid paying taxes and would buy additional shares at the current price for diversification reasons.

The second reason why the difference between the true and average tax basis policies is so small is that both the total amount of taxes paid by the investor and the total amount of tax rebates obtained by the investor are very small. In particular, we find that the tax-to-wealth ratio for the true tax basis policy is always between -2% and 2% , and is on average less than 1% for all time periods. Again, there is little disadvantage from approximating the true tax basis by the weighted average purchase price when the tax payments and rebates are such a small proportion of total wealth. Consequently, the stock-to-wealth ratio for the average tax basis policy is very similar to that of the true tax basis policy shown in the top panel of Figure 1.

4.2. Sensitivity to Parameter Values for Stock Returns and Risk Aversion

To understand whether the small difference in the exact and approximate tax basis policies is driven by our choice of benchmark parameter values, we analyze how the difference between these policies depends on our choice of parameter values for the expected annual rate of return and volatility of the stocks, and for the investor’s risk aversion. We solve problems with expected returns ranging from 8% to 14% , volatilities ranging from 15% to 25% , and relative risk-aversion parameters ranging from 2 to 4. We keep the rest of the parameters as in our base case, and we report only the comparison based on the certainty equivalent measure.

The certainty equivalents associated with the optimal policy and the four suboptimal policies are given in Table 1. The first column of this table gives the investor’s risk-aversion parameter (γ), the second column gives the expected annual rate of return (μ) on the stock, and the third column gives the volatility

Table 1 Sensitivity to Risk Aversion and the Mean and Volatility of Stock Returns

Parameters			CEQ		CEQ loss		
γ	μ	σ	(trubase)	(avgbase)	(buyhold)	(realize)	(augbuy)
2	0.08	0.15	1.8206	0.00%	0.39%	1.42%	0.02%
		0.20	1.7859	0.00%	0.38%	0.87%	0.05%
		0.25	1.7697	0.00%	0.36%	0.60%	0.08%
	0.10	0.15	2.0336	0.00%	0.66%	2.80%	0.05%
		0.20	1.9003	0.00%	0.67%	1.72%	0.06%
		0.25	1.8411	0.00%	0.69%	1.15%	0.16%
	0.12	0.15	2.4459	0.07%	1.42%	4.15%	0.73%
		0.20	2.1038	0.00%	0.88%	2.67%	0.04%
		0.25	1.9631	0.00%	0.93%	1.78%	0.15%
	0.14	0.15	3.2259	0.32%	4.40%	5.48%	4.01%
		0.20	2.4322	0.02%	1.24%	3.61%	0.24%
		0.25	2.1496	0.00%	1.07%	2.49%	0.08%
3	0.08	0.15	1.5539	0.00%	0.26%	0.92%	0.02%
		0.20	1.5343	0.00%	0.26%	0.57%	0.05%
		0.25	1.5252	0.00%	0.25%	0.39%	0.06%
	0.10	0.15	1.6716	0.00%	0.42%	1.80%	0.02%
		0.20	1.5982	0.00%	0.48%	1.09%	0.08%
		0.25	1.5654	0.00%	0.50%	0.73%	0.16%
	0.12	0.15	1.8875	0.00%	0.51%	2.73%	0.04%
		0.20	1.7084	0.00%	0.62%	1.66%	0.07%
		0.25	1.6327	0.00%	0.72%	1.10%	0.23%
	0.14	0.15	2.2562	0.05%	0.76%	3.56%	0.49%
		0.20	1.8794	0.00%	0.69%	2.31%	0.03%
		0.25	1.7323	0.00%	0.85%	1.51%	0.22%
4	0.08	0.15	1.4703	0.00%	0.20%	0.68%	0.02%
		0.20	1.4565	0.00%	0.20%	0.42%	0.04%
		0.25	1.4500	0.00%	0.19%	0.29%	0.05%
	0.10	0.15	1.5524	0.00%	0.33%	1.31%	0.02%
		0.20	1.5014	0.00%	0.38%	0.79%	0.09%
		0.25	1.4784	0.00%	0.40%	0.53%	0.15%
	0.12	0.15	1.6990	0.01%	0.35%	2.00%	0.01%
		0.20	1.5775	0.00%	0.51%	1.19%	0.11%
		0.25	1.5254	0.00%	0.60%	0.79%	0.24%
	0.14	0.15	1.9381	0.00%	0.26%	2.63%	0.05%
		0.20	1.6927	0.00%	0.55%	1.63%	0.07%
		0.25	1.5939	0.00%	0.74%	1.07%	0.28%

Note. This table shows the certainty equivalent (CEQ) value for different values of the investor’s risk aversion (γ) and the parameters governing the mean (μ) and volatility (σ) of stock returns, which are given in the first three columns. The fourth column gives the certainty equivalent of the true basis policy (trubase) and the last four columns give the percentage loss in certainty equivalent incurred by each of the four suboptimal policies: the average-basis policy (avgbase), the buy-and-hold policy (buyhold), the realize-all-capital-gains-and-losses policy (realize), and the augmented buy-and-hold policy (augbuy).

of stock returns (σ). The fourth column gives the certainty equivalent of the true basis policy (trubase), and the last four columns give the percentage loss in certainty equivalent incurred by each of the four suboptimal policies: the average-basis policy (avgbase), the buy-and-hold policy (buyhold), the realize-all-capital-gains-and-losses policy (realize), and the augmented-buy-and-hold policy (augbuy).

The main conclusion from this table is that the difference between the exact and approximate policies is larger when the expected stock return is larger,

the volatility is smaller, and the risk-aversion parameter is smaller. All these cases correspond to the situation where an investor holds a larger proportion of her wealth in stock, and consequently, the effect of capital gains tax is greater. Note that the largest percentage loss in using the average tax basis policy is 0.32%, and occurs for the portfolio problem with risk aversion 2, stock expected rate of return of 14%, and volatility of 15%. Herein, we refer to these parameter values as the “*stock-only case*” because it involves holding a levered position in stock; the stock-only case is not very realistic because it leads to a stock-to-wealth ratio larger than 2 for all dates and states, but this extreme case is useful for obtaining an upper bound to the certainty equivalent loss that an investor may incur when using the average tax basis policy.

4.3. Sensitivity to Model Refinements


Given that even for extreme parameter values the difference between the optimal portfolio policy and the average tax basis policy is quite small for the basic model, in this section we analyze how the portfolio policies are affected by a variety of refinements to the basic model, such as the introduction of transaction costs, stock dividends, intermediate consumption, labor income, tax reset at the last date, asymmetric taxation of long-term and short-term capital gains, and wash-sale constraints. A detailed analysis of each of these model refinements is given below in §§4.3.1–4.3.7, with a summary of the results presented in §4.3.8.

4.3.1. Transaction Costs. We examine the portfolio policies under the true and average tax basis for the base case problem in the presence of a 1% proportional transaction cost. Figure 2 gives the stock-to-wealth ratio of the true tax basis policy.

Not surprisingly, Figure 2 shows that in the presence of transactions costs the investor holds a portfolio that is much less diversified relative to the base case without transactions costs. In particular, the investor’s stock-to-wealth ratio ranges from 30% to 65% in the presence of transaction costs, whereas it ranges from 53% to 66% in the absence of transaction costs.

Moreover, by examining the computed true tax basis policy we note that transactions costs also deter the investor from using wash sales. Consequently, the investor holds shares with two or more different tax basis more frequently. In particular, we note that the investor holds up to 25% of stock bought at secondary tax basis for some of the paths for which the stock price decreases at $t = 2$. Thus, one might expect the difference in certainty equivalent under the true and exact policies to be larger. However, a second effect of the transaction costs is that it reduces trading volume so that the investor buys and sells less. In particular,

Figure 2 True Tax Basis Policy for the Base-Case Problem with Transaction Costs



Note. This figure graphs the stock-to-wealth ratio of the true tax basis policy for the base-case problem with a proportional transaction cost of 1%. The stock-to-wealth ratio is given for each of the 128 paths in the binomial tree and for each of the first seven dates.

we observe that, when following the true tax basis policy, the investor starts out holding 0.44 shares at date $t = 0$ and keeps the same 0.44 shares for almost all paths for which the stock price goes up at $t = 2$. The reduced trading volume offsets the benefits from using the true tax basis policy in the presence of multiple tax basis. As a result, the difference between the certainty equivalent of the exact and average tax policies remains negligible for the benchmark parameter values (below 0.01%) and is reduced in the presence of transaction costs from 0.32% to 0.06% for the stock-only parameter values.

4.3.2. Dividends. We compute the true and average tax basis policies for an investment problem with a stock with a nominal dividend yield of $d = 1.85\%$ and an expected before-dividend rate of return of 8%. Thus, the total rate of return (including dividends) is 10% as in our base case. We keep the rest of the parameters equal to their benchmark values given in §4.1.

An interesting effect of dividends is that they make the stock less attractive to the investor. The reason is that dividends are taxed and thus the investor loses part of her option to defer gains. As a result, the investor holds only 51% of her wealth in the stock in the presence of dividends while she holds 53% in our base case.

Our numerical results show that the difference in certainty equivalent between the true and average tax basis is very small also in the presence of dividends. In particular, the difference remains negligible for the benchmark parameters and remains at 0.32% for the stock-only parameter values.

4.3.3. Intermediate Consumption. If an investor needs to finance intermediate consumption, then she will need to sell stock and realize capital gains. Thus, one might expect that the desire for intermediate consumption will lead to an increase in the certainty equivalent difference under the exact tax basis and average tax basis policies.

We solve our base-case problem for an investor that derives utility from consumption at all dates. We observe that when intermediate consumption is introduced, the investor buys stock at the first date and then sells stock at every subsequent date to finance consumption. Consequently, all stock held by the investor has the same tax basis. This explains why the true and average tax basis policies yield the same certainty equivalents both for the benchmark parameter values and the stock-only parameter values; there is no advantage in using the true tax basis if all the stock held by the investor has the same tax basis.


4.3.4. Labor Income. In the presence of labor income, the investor will need to reinvest a part of her nonfinancial income to have a diversified portfolio, which will lead to stock holdings with multiple tax basis.

We solve our base case problem with the additional assumption that, at each date, the investor obtains 15% of her pre-tax wealth as labor income.¹⁵ Figure 3 gives the stock-to-wealth ratio for the true tax basis policy.

By examining the true tax basis policy, we note that when the investor obtains labor income, she buys stock more frequently. In particular, the investor sometimes buys stock even when the stock price *increases*. As a consequence of this effect, the investor holds stock with more than one tax basis more often than in the absence of labor income. But another effect of labor income is that the investor rarely sells any stock. These two effects offset each other and the difference in certainty equivalent between both policies changes very little in the presence of labor income. In particular, it remains very small for the base-case parameter values (0.050%), and decreases only slightly (from 0.31% to 0.29%) for the stock-only parameter values.

4.3.5. Tax Reset Provision at the Last Date. Dammon et al. (2001) show that the effect of tax forgiveness on the investor's portfolio policies depends significantly on the investor's *bequest* motives; that is, on how much the investor values his own consumption as compared to that of her beneficiaries' after her death. To analyze this dependence we consider two different cases: (i) a case with intermediate and terminal consumption, where the bequest motive is

Figure 3 True Tax Basis Policy for the Base-Case Problem with Labor Income



Note. This figure graphs the stock-to-wealth ratio of the true tax basis policy for the base-case problem with labor income of 15% of the pre-tax wealth. The stock-to-wealth ratio is given for each of the 128 paths in the binomial tree and for each of the first seven dates.

represented by the consumption at the terminal date, and (ii) a model without intermediate consumption, where the investor has only bequest motives. That is, the case with intermediate consumption represents a situation with weaker bequest motives.

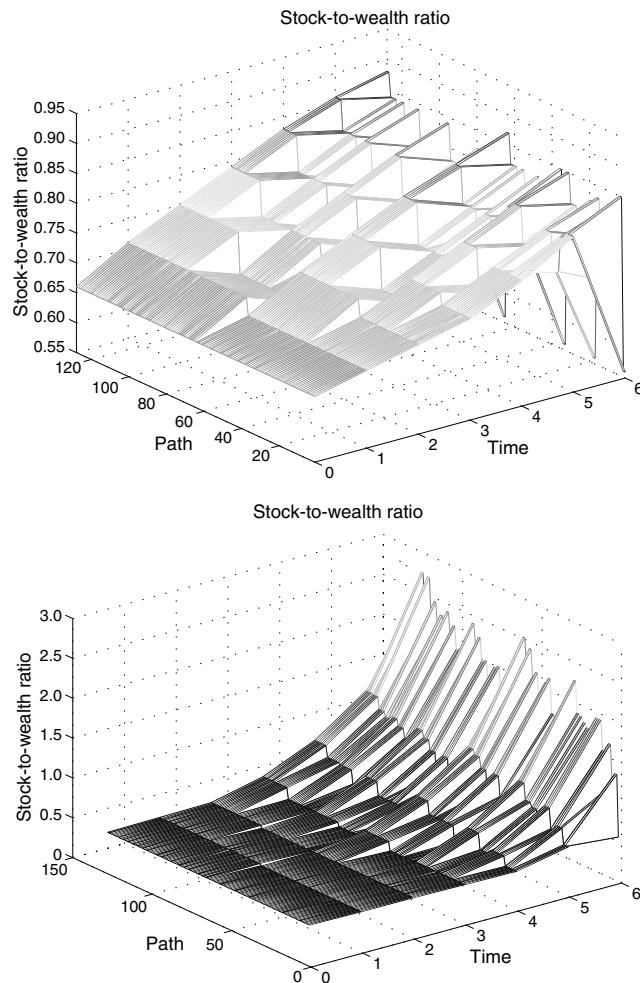
We compute the true and average tax basis policies with and without intermediate consumption in the presence of the tax reset provision at death. The top and bottom panels of Figure 4 give the stock-to-wealth ratio for the cases without and with intermediate consumption, respectively.

We find that tax forgiveness makes the stock a more attractive asset. In particular, the investor's stock-to-wealth ratio is higher. For instance, in the case without intermediate consumption and tax forgiveness the ratio ranged between 56% and 87%, whereas in the base case without intermediate consumption it ranged between 53% and 66%. In the case with intermediate consumption and tax forgiveness the ratio ranged between 50% and 273%, whereas in the base case with intermediate consumption it ranged between 54% and 65%.

Moreover, Dammon et al. (2001) show that in the presence of the tax reset provision, an investor tends to hold a larger proportion of stock as she becomes older. Figure 4 shows this is true even for the optimal portfolio policy obtained using the true tax basis. In particular, the stock-to-wealth ratio increases with time even for the first 64 paths, where the stock price goes down at $t = 2$. In addition, note that this effect is more pronounced in the case with intermediate consumption because the investor prefers to borrow cash to finance intermediate consumption and holds on to stocks to take full advantage of tax forgiveness.

¹⁵ Dammon et al. (2001) model labor income in the same manner.

Figure 4 True Tax Basis Policy for the Base-Case Problem with Tax Reset Provision at Death



Note. This figure graphs the true tax basis policy for the base-case problem with the tax reset provision at death. The top panel gives the stock-to-wealth ratio for the case without intermediate consumption, for each of the 128 paths in the binomial tree and for the first seven dates. The bottom panel gives the stock-to-wealth ratio for the case with intermediate consumption.

As a result, the stock-to-wealth ratio grows large in the last few periods. This insight matches the one in Dammon et al. (2001) that, with a weaker bequest motive, the growth of the stock-to-wealth ratio with investor's age is more pronounced.

Also, Figure 4 shows that at the penultimate date the investor sells all those shares whose tax basis is higher than the current stock price to capture the tax rebate while continuing to hold all those shares for which the tax basis is below the current stock price to take advantage of the tax reset provision at death.

Finally, the difference in certainty equivalent between the true and average tax basis policies remains very small even in the presence of the tax reset provision at death. In particular, the difference is zero for the case with intermediate consumption both with the base-case and the stock-only parameter values

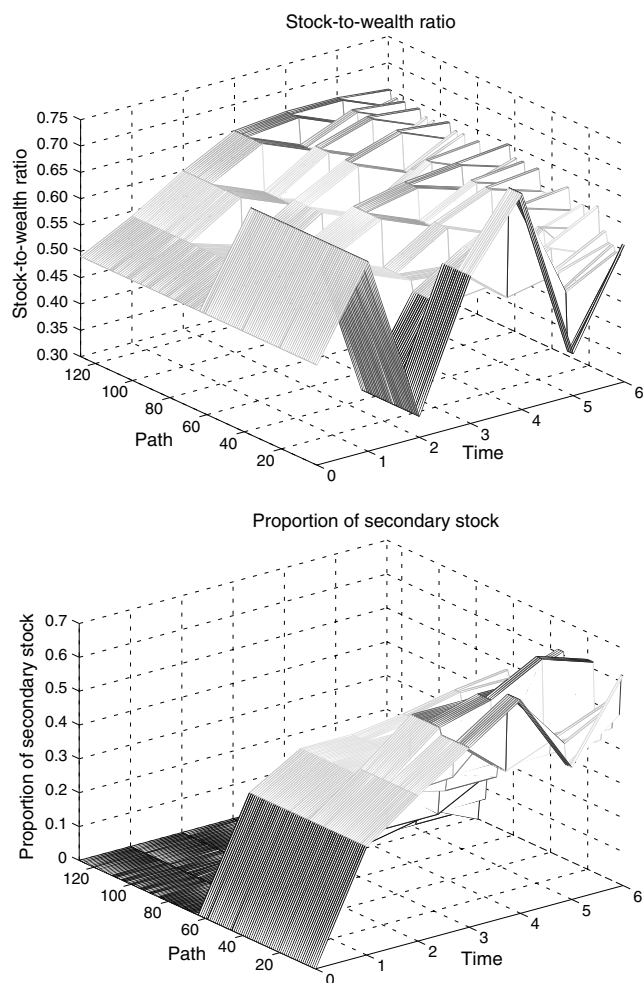
because the investor buys stock only on the first date. For the case without intermediate consumption, the difference remains very small for the base-case parameter values (below 0.01%), but increases from 0.31% to 1.46% for the stock-only parameter values. The explanation for this larger certainty equivalent loss is twofold. First, with tax forgiveness at death the investor holds more stock in her portfolio. Second, with tax forgiveness at death the investor benefits more from selling those shares with the highest tax basis first because she keeps the shares with the lowest tax basis until the final date, and then she does not have to pay any taxes on them. Thus, the investor may avoid paying taxes all together on those shares that she bought at the lowest prices.

4.3.6. Asymmetric Taxation. In this section, we discuss how the existence of different short-term tax rates (for shares bought less than a year ago) and long-term tax rates (for shares bought more than a year ago) affects the difference between the true and average tax basis policies. Constantinides (1984) shows that the existence of a lower long-term tax rate may persuade investors to realize all capital gains at the long-term tax rate simply by realizing the gain one day after the end of the year to reset the option to realize short-term losses the following year (by realizing any losses one day before the end of the following year). In particular, Constantinides (1984) gives conditions under which it is optimal to realize all capital gains and losses every year.

Our base-case and stock-only case (assuming a short-term tax rate of 35% and a long-term tax rate of 20%) satisfy the conditions given by Constantinides (1984). Thus, in the presence of asymmetric long- and short-term tax rates, the true and average tax basis policies are exactly equal for both the base-case and the stock-only parameter values.

4.3.7. Wash-Sale Constraints. In the presence of wash-sale constraints, an investor with an embedded capital loss has to choose between buying stock to diversify her portfolio or selling stock to obtain a tax rebate because the constraint rules out undertaking both activities simultaneously. As a consequence, the investor with an embedded capital loss is more likely to hold stock with more than one tax basis than the investor that is not wash-sale constrained.

We solve our problem for the base-case and stock-only parameters in the presence of wash-sale constraints. Figure 5 gives the stock-to-wealth ratio and the proportion of secondary stock, respectively. The lower panel in Figure 5 shows that in the presence of a constraint on wash sales, the investor holds stock with more than one tax basis for the first 64 paths, where the price goes down after the first date. As a consequence, the difference between using the true or

Figure 5 True Tax Basis Policy for the Base-Case Problem with Wash-Sale Constraints

Note. This figure graphs the true tax basis policy for the base-case problem with wash-sale constraints. The top panel gives the stock-to-wealth ratio for each of the 128 paths in the binomial tree and for each of the first seven dates. The bottom panel gives the proportion of secondary stock; that is, the proportion of stock held at a tax basis different from the basis of the majority of the stock.

the average tax basis policies is expected to be larger than in the absence of wash-sale constraints. In fact, the difference in certainty equivalent between both policies goes up to 0.047% for the base-case parameter values. However, the difference for the stock-only parameter values stays at 0.31%. The reason for this is that for the parameter values in the stock-only case, when the stock price goes down, it goes down by a factor very close to one (0.99). As a consequence, the effect of wash sales has little importance for the stock-only parameter values.

Finally, note that the first panel in Figure 5 also confirms the intuition in Garlappi et al. (2001) that a wash-sale constrained investor with a capital loss buys more *and* sells more than an investor who is not wash-sale constrained. In particular, note how

Table 2 Sensitivity to Model Refinements

Model refinement	Base-case parameters		Stock-only parameters	
	CEQ (trubase)	CEQ loss (avgbase)	CEQ (trubase)	CEQ loss (avgbase)
None: Basic model	1.5982	0.00%	3.2258	0.31%
Transaction costs	1.5642	0.00%	2.9374	0.06%
Stock dividends	1.6945	0.00%	4.7614	0.57%
Intermediate consumption	0.0554	0.00%	0.0269	0.00%
Labor income	4.4249	0.05%	10.165	0.29%
Tax reset at last date	1.7467	0.00%	4.8011	1.47%
Asymmetric taxation	—	0.00%	—	0.00%
Wash-sale constraints	1.5946	0.05%	3.2187	0.31%

Note. This table gives the certainty equivalent (CEQ) of the true and average tax basis policies first for the basic model and then for several model refinements for both the base-case parameter values and the stock-only parameter values.

the investor's stock-to-wealth ratio oscillates between high and low values for the first 64 paths in the first panel of Figure 5.

4.3.8. Summary of CEQ Loss Under Various Model Refinements. In this section, we summarize the results from the analysis of the various model refinements considered above. We report in Table 2 how the difference between the certainty equivalent of the true and average tax basis policies is affected by the various model refinements considered below. The first column of the table describes the particular refinement being considered. The second column gives the certainty equivalent of the true basis policy (trubase) and the third column gives the percentage loss in certainty equivalent incurred by the average-basis policy (avgbase) under the base-case parameters, while the fourth and fifth columns report the same quantities for the stock-only parameters. The results for the stock-only parameters effectively gives an upper bound to the certainty equivalent loss incurred by an investor using the average tax basis policy.

From Table 2, we see that for the more realistic base-case parameter values, the certainty equivalent (CEQ) loss from using the suboptimal policy based on the average tax basis is always less than 0.1%. Even for the more extreme stock-only parameter values, the CEQ loss from using a suboptimal portfolio strategy based on the average tax basis is always less than 1.5%.

4.4. Sensitivity to Number of Stocks and Correlation

So far, we considered the portfolio problem of an investor who can invest in only a single stock in addition to the risk-free asset. The portfolio problem with multiple risky assets in the presence of capital gains tax has been studied in detail in Dammon et al. (2002), Garlappi et al. (2001), and Gallmeyer et al. (2001). In contrast to these papers, where the tax basis is

Table 3 Sensitivity to Number of Stocks and Correlation

Correlation	Base-case parameters		Stock-only parameters	
	CEQ (trubase)	CEQ loss (avgbase)	CEQ (trubase)	CEQ loss (avgbase)
0.5	1.63074	0.00%	2.76534	0.25%
0.9	1.59912	0.00%	2.20368	0.09%

Note. This table gives the certainty equivalent (CEQ) of the true tax basis policy (trubase) and an upper bound on the percentage loss in certainty equivalent incurred by the average tax basis policy (avgbase) for problems with seven periods and two stocks and for correlation coefficient between the returns of the stocks equal to $\rho = \{0.5, 0.9\}$ for both the base-case parameter values and stock-only parameter values.

approximated using the weighted average purchase price, we consider the portfolio problem using the exact tax basis.

We consider the problem with seven periods when there are two stocks whose rates of return are correlated with correlation coefficients equal to 0.5 or 0.9. All other parameter values are the same as in the case with a single stock. In particular, the stocks have an expected annual rate of return of 10%, annual volatility of 20%, and it is assumed that they pay no dividends. We compute the true tax basis policy using both the benchmark parameter values and stock-only parameter values for the two stocks. We also compute an upper bound on the certainty equivalent loss incurred by using the average tax basis policy.¹⁶

Table 3 gives the certainty equivalent of the true tax basis policy and the upper bound on the loss in certainty equivalent incurred by the average tax basis policy. We find that the difference in certainty equivalent between the true and average tax basis policies is negligible (below 0.01%) for the base-case parameter values and small (below 0.25%) for the stock-only parameter values. Moreover, the insight that the investor rarely holds stocks with more than one different tax basis still holds when there are two risky assets. In particular, the investor holds less than 1% of secondary stock on average for the base-case parameters and for correlation coefficients of 0.5 and 0.9. As expected, the certainty equivalent of the true tax basis policy decreases when the correlation coefficient increases.

4.5. Sensitivity to Number of Periods

In our analysis so far, we have considered problems with only seven periods. We now solve a sequence of problems with the number of periods $T = \{7, 8, 9, 10\}$

¹⁶ This upper bound is generated by computing the certainty equivalent obtained by an investor who holds the same number of shares as the investor following the true tax basis policy, but uses the weighted average purchase price as the tax basis. This is a lower bound on the certainty equivalent obtained by using the average tax basis policy, and thus provides an upper bound on the certainty equivalent loss.

Table 4 Sensitivity to Number of Periods

Number of periods (T)	Base-case parameters		Stock-only parameters	
	CEQ (trubase)	CEQ loss (avgbase)	CEQ (trubase)	CEQ loss (avgbase)
7	1.59819	0.000%	3.22585	0.32%
8	1.71128	0.000%	3.84603	0.45%
9	1.83283	0.000%	4.59182	0.56%
10	1.96346	0.004%	5.48725	0.78%

Note. This table gives the certainty equivalent of the true tax basis policy (trubase) and the percentage loss in certainty equivalent incurred by the average tax basis policy (avgbase) for problems with number of periods equal to $T = \{7, 8, 9, 10\}$ for both the base-case parameter values and stock-only parameter values.

for both the base-case parameter values and stock-only parameter values.

Table 4 gives the certainty equivalent for the true and average basis policies. Note that the difference in certainty equivalent between the true and average basis policies remains negligible for problems of up to 10 periods under the base-case parameter values. On the other hand, for the less realistic stock-only parameter values, the difference in certainty equivalent grows with the number of periods, but stays below 1% for problems with a horizon of $T = 10$.

5. Conclusions

We have shown how to compute the optimal consumption and portfolio policies of an investor subject to capital gains taxes using the *exact* tax basis rather than the approximation used in the literature, the weighted average purchase price. This is made possible by recognizing the favorable features of the problem—concavity and separability of the objective function, linearity of the constraints, and sparsity of the objective function and constraints—and using nonlinear programming rather than dynamic programming.

We determine the optimal portfolio policies for problems of up to seven periods and two stocks, or 10 periods and one stock. The ability to solve larger scale problems using the exact tax basis allows us to compare the optimal policy with the suboptimal policy obtained by approximating the tax basis by the weighted average purchase price. We find that the certainty equivalent loss from using the suboptimal policy based on the weighted average purchase price is less than 1%. This certainty equivalent loss is larger when the expected stock return is larger, the volatility is smaller, and the risk-aversion parameter is smaller, but it is still less than 1% even for relative risk aversion ranging from two to four, for expected stock returns ranging from 8% to 14%, and for the volatility of stock returns ranging from 15% to 25%. This is also true in the presence of transaction costs,

dividends, intermediate consumption, labor income, tax reset provision at death, wash-sale constraints, and two risky stocks instead of one.

Our analysis also allows us to get new insights about the properties of the optimal portfolio policy and to confirm some of the findings in the literature based on approximating the true tax basis by the weighted average purchase price. For example, we find that an investor following the optimal investment policy rarely holds shares bought at more than one date: the proportion of stock held at a basis different to the tax basis of most of the stock is typically less than 10%, and is lower than 1% on average. In the presence of a capital gains tax, the investor also reduces the volume of trading; consequently, the investor's stock-to-wealth ratio varies over time whereas in the absence of taxes it would be constant. Introduction of transaction costs lead the investor to reduce trading further and also deter the investor from using wash sales. The presence of cash dividends and labor income, on the other hand, leads to constant rebalancing of the investor's portfolio, and hence, an increase in the proportion of stock with multiple tax bases. As in Dammon et al. (2001), we find that in the presence of the tax reset provision an investor tends to hold a larger proportion of stock as she becomes older, and similar to the result in Garlappi et al. (2001), we find that a wash-sale constrained investor with a capital loss buys more and sells more than an investor who is not wash-sale constrained.

Last but not least, just like the work of Dybvig and Koo (1996) inspired us to find a better way for solving the problem of portfolio selection in the presence of taxes on capital gains, we hope that this paper will lead other researchers to find new ways for attacking this challenging problem.

Acknowledgments

The authors thank an anonymous referee and Suleyman Basak for several suggestions. We gratefully acknowledge comments from Viral Acharya, Joao Cocco, Zeger Degraeve, James Dow, Phil Dybvig, Julian Franks, Lorenzo Garlappi, David Goldreich, Francisco Gomes, Hyeng Keun Koo, Francisco Nogales, Stathis Tompaidis, Stavros Zenios, Harold Zhang, and seminar participants at London Business School, the Judge Institute of Management, the 2003 annual meeting of INFORMS in Atlanta, and the 2004 meetings of the Western Finance Association in Vancouver.

References

- Basak, S., B. Croitoru. 2001. Nonlinear taxation, tax arbitrage and equilibrium asset prices. *J. Math. Econom.* **35** 347–382.
- Basak, S., M. Gallmeyer. 2003. Capital market equilibrium with differential taxation. *Eur. Finance Rev.* **7** 121–159.
- Birge, J., F. Louveaux. 1997. *Introduction to Stochastic Programming*. Springer-Verlag, New York.
- Clarke, F. 1983. *Optimization and Nonsmooth Analysis*. John Wiley & Sons, New York.
- Constantinides, G. 1983. Capital market equilibrium with personal taxes. *Econometrica* **51** 611–636.
- Constantinides, G. 1984. Optimal stock trading with personal taxes: Implications for prices and the abnormal January returns. *J. Financial Econom.* **13** 65–89.
- Dammon, R., R. Green. 1987. Tax arbitrage and the existence of equilibrium prices for financial assets. *J. Finance* **42** 1143–1166.
- Dammon, R., C. Spatt, H. Zhang. 2001. Optimal consumption and investment with capital gains taxes. *Rev. Financial Stud.* **14** 583–616.
- Dammon, R., C. Spatt, H. Zhang. 2002. Diversification and capital gains taxes with multiple risky assets. Working paper, Carnegie Mellon University, Pittsburgh, PA.
- Dammon, R., C. Spatt, H. Zhang. 2004. Optimal asset location and allocation with taxable and tax-deferred investing. *J. Finance* **59** 999–1037.
- DeMiguel, A., F. Nogales. 2002. On the relationship between bilevel decomposition algorithms and Schur interior-point methods. Working paper, London Business School, London, U.K.
- Dybvig, P., H. Koo. 1996. Investment with taxes. Working paper, Washington University, Saint Louis, MO.
- Gallmeyer, M., R. Kaniel, S. Tompaidis. 2001. Tax management strategies with multiple risky assets. Working paper, University of Texas, Austin, TX.
- Garlappi, L., V. Naik, J. Slive. 2001. Portfolio selection with multiple assets and capital gains taxes. Working paper, University of British Columbia, Canada.
- Gill, P., W. Murray, M. Saunders. 2002. SNOPT: An SQP algorithm for large-scale constrained optimization. *SIAM J. Optim.* **12** 979–1006.
- Gupta, A., W. Murray. 2000. Optimal investment with behavioral utilities using a binomial tree model for asset-returns. Discussion paper SOL 00-1, Department of Management Science and Engineering, Stanford University, Stanford, CA.
- He, H. 1990. Convergence from discrete to continuous-time contingent claims prices. *Rev. Financial Stud.* **3** 523–546.
- Luenberger, D. 1984. *Linear and Nonlinear Programming*. Addison-Wesley, Reading, MA.
- Luo, Z.-Q., J.-S. Pang, D. Ralph. 1996. *Mathematical Programs with Equilibrium Constraints*. Cambridge University Press, Cambridge, U.K.
- Nocedal, J., S. Wright. 1999. *Numerical Optimization*. Springer-Verlag, New York.
- Schaefer, S. 1982. Taxes and security market equilibrium. W. Sharpe, C. Cootner, eds. *Financial Economics: Essays in Honor of Paul H. Cootner*. Prentice-Hall, NJ.