

The Optimal Trading and Pricing of Securities with Asymmetric Capital Gains Taxes and Transaction Costs

Robert M. Dammon
Chester S. Spatt
Carnegie Mellon University

This article explores the optimal trading and pricing of taxable securities with asymmetric capital gains taxes and transaction costs. In the long-term region, investors realize all gains below some critical cutoff level, which we derive analytically. In the short-term region, investors defer all gains and, depending upon the time remaining in the short-term region, may also defer small losses. Contrary to common intuition, deferral of short-term losses can be optimal even without transaction costs. The value of tax timing is considerably higher under the optimal trading strategy than under alternative strategies previously analyzed. The impact of offset rules is also explored.

The effect of capital gains taxes on the optimal trading and equilibrium pricing of taxable securities is an

We gratefully acknowledge valuable comments from Peter Bossaerts, Steve Heston, Franklin Allen (the editor), two anonymous referees, and the seminar participants at Carnegie Mellon University, Columbia University, Cornell University, Duke University, Hebrew University of Jerusalem, the University of Maryland, the University of Minnesota, Northwestern University, the University of Pennsylvania, Princeton University, the University of Southern California, Tel-Aviv University, the Third International Finance Conference of the Groupe Haute Etudes Commerciales, Toulouse University, and the Western Finance Association Meetings. We also gratefully acknowledge the research assistance of Bernt Oedegaard and Cliff Parsons and financial support from the National Science Foundation. Previous versions of this article were circulated under the title "Optimal Trading Policies and the Pricing of Securities with Asymmetric Taxation." Address correspondence to the authors at the Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh, PA 15213.

important theoretical and empirical issue. In this article we focus on two important aspects of capital gains taxation that influence investors' trading behavior. First, under current U.S. tax law, gains and losses on capital assets are not taxed until the investor sells the asset. Second, the tax rate that applies to capital gains and losses may be a function of the investor's holding period. For example, under current U.S. tax law, capital gains and losses on assets held longer than 1 year (referred to as long-term) are taxed at a maximum rate of 28%, whereas capital gains and losses on assets held less than 1 year (referred to as short-term) are taxed at a maximum rate of 39.6%.¹ These features give investors the incentive to time their asset sales so as to minimize the tax burden of owning taxable securities. The purpose of this article is to derive the optimal trading strategy for investors in the presence of capital gains taxes and transaction costs, and to explore the implications of the optimal trading strategy for asset pricing.²

Under asymmetric taxation (i.e., when the long-term tax rate is less than the short-term tax rate), Constantinides (1984) has shown that it can be optimal to sell and repurchase an asset with an embedded long-term capital gain to reset the tax basis and restart the option to realize future capital losses short term. While this is an important theoretical insight, it does not offer explicit guidance as to the actual trading policy that is optimal in the presence of asymmetric taxation. What is needed is an optimal realization boundary in both the short-term and long-term regions above which all capital gains (and losses) are deferred and below which all capital gains (and losses) are realized. Using a simple discrete-time model, we derive the optimal realization boundary and explore its qualitative properties. In the long-term region, the optimal realization boundary is time independent and can be derived analytically. We explore how the optimal long-term realization boundary is affected by the level of transaction costs, the volatility of the stock, the long-term and short-term tax rates, and the length of the short-term region.

Because our model allows for multiple trading dates and stock price changes within the short-term region, a number of interesting and counterintuitive results emerge. We show that the optimal trading policy in the short-term region involves the deferral of all capital gains and, depending upon the time remaining in the short-term region, can also involve the deferral of small capital losses, even in the absence

¹ In the U.S., the ratio of the long-term tax rate to the short-term tax rate has varied over time, ranging between 40% and 100% during most of the postwar period.

² Recently, Bossaerts and Dammon (1994) estimated and tested a consumption-based asset pricing model that incorporates tax-timing options and provided some empirical evidence supporting the existence of capital gains tax effects in asset prices.

of transaction costs. This result is in direct opposition to the common belief that investors should optimally realize short-term capital losses as soon as they occur in the absence of transaction costs. The intuition underlying our result stems from the fact that in addition to the tax benefits of realizing short-term losses there is also a cost associated with restarting the short-term holding period *prior* to reaching the long-term region.

To illustrate, consider an investor who has held an asset for a period of 6 months, where the holding period for long-term tax treatment is 1 year. Assume that the market price of the asset is equal to the investor's tax basis so that currently there is neither a gain nor a loss on the position. Although the investor would pay no tax if he sold and repurchased the asset at this time, doing so would make the investor worse off by restarting the short-term holding period and increasing the length of time he must wait to qualify for preferential long-term tax treatment on any subsequent capital gains. In essence, with less time remaining in the short-term region the investor is better able to exploit the asymmetry between long-term and short-term tax rates. Thus, restarting the short-term holding period prior to reaching the long-term region has a cost that investors must take into account when deciding whether to realize short-term capital losses. We show that it can be optimal for investors to forego the tax benefits associated with realizing small short-term losses to avoid restarting the short-term holding period prematurely. The maximum short-term loss investors are willing to forego is a function of the time remaining until the asset qualifies for long-term tax treatment. Under plausible parameter values, we find that it can be optimal for investors to defer realizing short-term capital losses of 10% or more in the absence of transaction costs.

We assume a simple binomial dividend process for the taxable security and derive its equilibrium price endogenously to reflect the optimal tax-trading policy and transaction costs. We also assume that each taxable security has its own "tax-exempt counterpart," which can be traded without taxes or transaction costs. The price differential between the taxable and tax-exempt securities provides an estimate of the value of the tax-timing option (net of transaction costs) available on the taxable security. We explore the qualitative properties of the value of the tax-timing option using numerical methods. We find that for high- and medium-variance stocks, the tax-timing option accounts for a substantial fraction of the total market value of the taxable security. In the presence of offset rules, which require investors to net short-term losses against long-term gains realized within the same tax year, a lower bound on the value of tax timing can be obtained by assuming that investors never realize long-term capital gains. We find

that even when investors ignore the opportunity to realize long-term capital gains to restart the short-term holding period, the tax-timing option retains considerable value due to the ability to realize short-term capital losses. Thus, the ability to trade optimally within the short-term region is of critical importance.

This article is related to a number of theoretical and empirical studies in the literature. Constantinides (1984) analyzes the optimal trading of taxable securities in the presence of asymmetric taxation in a setting similar to ours, but where stock price changes and trading are assumed to occur only once per year. As a result, investors in his model *never* defer short-term capital losses and either realize *all* long-term capital gains each year or defer *all* long-term capital gains indefinitely.³ In contrast, our model allows for multiple stock price changes and trading dates within the year, which substantially alters investors' optimal realization behavior. In particular, we show that there exists an optimal long-term realization boundary above which all long-term gains (and losses) are deferred and below which all long-term gains (and losses) are realized. Unlike the Constantinides (1984) model, the optimal long-term realization boundary in our model can take on values between zero (i.e., always defer) and infinity (i.e., always realize). Moreover, we show that it can be optimal for investors to defer small short-term capital losses to avoid restarting the short-term holding period prematurely. Thus, the two models produce substantially different predictions regarding the effect of taxes on investors' optimal trading behavior and the pricing of financial assets.

While our model can be interpreted as a discrete-time version of the model developed in Williams (1985), he focuses on the optimal trading and equilibrium pricing of depreciable assets in a setting with symmetric capital gains taxes. In contrast, we examine the optimal trading and equilibrium pricing of financial assets with asymmetric capital gains taxes. Schultz (1988) analyzes a model with asymmetric capital gains taxation and transaction costs and uses numerical techniques to solve for the optimal realization boundary in the short-term region, but assumes that investors *never* realize long-term capital gains.⁴ In contrast, our characterization of the optimal realization pol-

³ In Constantinides's (1984) model, the stock price is assumed to follow a discrete binomial process with a single stock price change and trading date per year. At the end of the 1-year holding period, the investor has the option to realize or defer the gain or loss on the stock, and can elect either long-term or short-term tax treatment. Given this structure, investors optimally realize all losses short term and either realize all gains long term each year or defer all gains indefinitely. It is never optimal in Constantinides's (1984) framework to defer short-term losses.

⁴ If investors never realize long-term capital gains, then there is no benefit to deferring the realization of short-term capital losses in the absence of transaction costs. Consequently, in Schultz's (1988) model, investors realize all short-term capital losses as soon as they occur when transaction costs are zero. This is inconsistent with the optimal realization behavior derived in this article.

icy in the short-term region reflects the *optimal* long-term realization policy. Finally, Dammon, Dunn, and Spatt (1989) explore the empirical significance of tax-timing options using Monte Carlo techniques and conclude that the incremental value of the restarting option is generally much smaller than that found by Constantinides (1984).⁵ However, Dammon, Dunn, and Spatt (1989) do not explicitly solve for the optimal realization policy in the presence of asymmetric taxation, but examine similar heuristic trading policies to those analyzed by Constantinides (1984) in which trading is allowed only once per year. Our results indicate that by allowing for more frequent trading within the year, the value of the restarting option increases substantially. Consistent with Dammon, Dunn, and Spatt (1989), we also find that the value of restarting is sensitive to transaction costs and offset rules.

This article is organized as follows. Our assumptions and the basic model are discussed in Section 1. In Section 2, we derive the difference equation and associated boundary conditions that describe the investor's personal valuation of his holdings of the taxable security and discuss how the investor's personal valuation relates to the equilibrium market price. In Section 3, we analyze the investor's optimization problem and derive the optimal realization policy for the long-term region. We also discuss the dynamic programming approach that is used to solve for the optimal realization policy in the short-term region. In Section 4, we present some numerical examples to illustrate the optimal trading and equilibrium pricing of the taxable security. Qualitative properties of the optimal trading strategies and equilibrium asset pricing are also examined. Section 5 provides a comparison of the results for the optimal tax-trading strategy derived in this article with some alternative tax-trading strategies that have been examined previously in the literature. We also discuss the impact of offset rules on the value of the tax-timing option. Section 6 concludes.

1. The Model

The tax treatment of capital gains and losses in this article is a simplification of the actual U.S. tax code. Capital gains and losses are untaxed until the investor sells the asset. Realized capital gains and losses are taxed as short term if the asset has been held for no more than N trading periods and are taxed as long term otherwise. While the actual

⁵ The "restarting option" refers to the option to realize long-term capital gains to reset the tax basis and restart the short-term holding period. The value of the restarting option obviously depends upon the trade-off between the cost of realizing long-term capital gains and the value of the option to realize losses short term.

tax code requires the investor to offset gains and losses on different realizations to determine whether the investor has a net long-term or net short-term capital gain or loss, we ignore this complication of the tax code and assume that all realizations are taxed separately. The impact of offset rules on the value of the tax-timing option will be investigated later in this article. The tax rate on dividends is denoted by τ_D , the tax rate on long-term gains and losses is denoted by τ_L , and the tax rate on short-term gains and losses is denoted by τ_S . We assume that $\tau_L \leq \tau_S$, but leave the relationship between these tax rates and the dividend tax rate unspecified. Finally, we assume that there is no capital loss limitation or restrictions on wash sales.⁶

We assume an exogenous cash flow (i.e., dividend) process and derive the equilibrium price of the stock endogenously in conjunction with the optimal tax-trading policy. To provide for a richer set of possible realization strategies, we assume that there are N ($N \geq 1$) trading dates in the short-term region.⁷ At each trading date, the stock pays a dividend and the equilibrium ex-dividend price of the stock is determined. All trades take place at the ex-dividend price and all tax payments and rebates are made at the time they are incurred. Investors have infinite horizons and there are no forced asset sales.

The stock is assumed to generate pretax dividends per share of X_t dollars at date t , $X_t > 0$. To simplify the analysis, we assume that these dividends evolve through time according to a binomial process of the form:

$$X_t \begin{cases} X_{t+1} = uX_t, & \text{with probability } q \\ X_{t+1} = u^{-1}X_t, & \text{with probability } 1 - q \end{cases} \quad (1)$$

where u ($u > 1$) and q ($0 < q < 1$) are constant over time. There also exists a riskless, tax-exempt bond with a constant price of \$1 that yields a constant interest rate of $r \equiv R - 1$. Finally, we assume that there exists a tax-exempt security generating dividends according to the same binomial process as in Equation (1), but for which both dividends and capital gains and losses are tax exempt. We refer to this

⁶ The current U.S. tax code limits the investor's net losses in any given year to a maximum of \$3,000. Unused capital losses may be carried forward indefinitely, but the investor loses the time value of the tax deduction during the interim. The deferral of the tax benefits due to the capital loss limitation will reduce the value of the tax-timing option to some degree, but it will not eliminate entirely the value of tax timing. The wash-sale rule disallows the deduction of a capital loss on an asset that is repurchased within 30 days of its sale. This rule can be circumvented easily by using the proceeds of the sale to purchase a different asset with similar risk and return characteristics.

⁷ N is the number of trading dates (not including the date of purchase) that occur before the holding period becomes long term. If the asset is sold N trading periods after its purchase, the gain or loss is treated as short term. Long-term treatment of gains and losses does not begin until $N + 1$ trading periods after the date of purchase.

security as the stock's "tax-exempt counterpart" and assume that it is traded in a competitive capital market without transaction costs.⁸

We denote by π_u and π_d the equilibrium state prices for \$1 of after-tax income in the "up state" and "down state," respectively. For simplicity, we assume that investors are risk neutral. Under risk neutrality, the tax-exempt interest rate, $r \equiv R - 1$, is equal to investors' personal rate of discount for intertemporal utility. This produces constant equilibrium state prices given by $\pi_u = R^{-1}q$ and $\pi_d = R^{-1}(1 - q)$. Under these conditions, the equilibrium ex-dividend price of the tax-exempt security at date t , denoted $\hat{P}_t$, is given by⁹

$$\hat{P}_t = \hat{\Pi} X_t, \quad (2)$$

where

$$\hat{\Pi} = \frac{\pi_u u + \pi_d u^{-1}}{1 - \pi_u u - \pi_d u^{-1}}. \quad (3)$$

We refer to $\hat{\Pi}$ as the tax-exempt pricing operator and note that it is independent of time and equal to the reciprocal of the tax-exempt security's dividend yield. Consequently, in equilibrium, the ex-dividend price of the tax-exempt security evolves according to the same binomial process that governs the dividend.

The taxable security is traded in a competitive market with proportional transaction costs. We denote the ex-dividend price of the taxable security at date t by P_t . In the presence of taxes, the value that an investor places on his current holdings of the taxable security may differ from its market price. Let $W(P_t, B, b)$ denote the investor's personal valuation of a position in one share of the taxable security with current price P_t , tax basis B , and age (or holding period) b . In other words, the investor is indifferent at the margin between holding one share of stock with price P_t , basis B , and age b , or having $W(P_t, B, b)$ after-tax dollars. In equilibrium, both P_t and $W(P_t, B, b)$

⁸ The model could have incorporated transaction costs for the tax-exempt security without materially affecting the subsequent analysis.

⁹ Equations (2) and (3) do not require risk neutrality, but would emerge under any set of conditions that lead to an equilibrium in which state prices are constant through time. To prevent arbitrage, the equilibrium price of the tax-exempt security must be proportional to its dividend (i.e., doubling the dividend must double the price) and, with constant equilibrium state prices, the constant of proportionality will be independent of time. Thus, the equilibrium pricing relationship in Equation (2) follows directly from the absence of arbitrage and constant equilibrium state prices. Equation (3) can be derived from Equation (2) by first writing the date t price of the tax-exempt security as the discounted value of the one-period-ahead cash flows:

$$\hat{P}_t = \pi_u [u X_t + u \hat{P}_t] + \pi_d [u^{-1} X_t + u^{-1} \hat{P}_t],$$

where $u \hat{P}_t$ and $u^{-1} \hat{P}_t$ are the equilibrium prices of the tax-exempt security at date $t + 1$ in the "up state" and "down state," respectively. Substituting Equation (2) into the above equation and solving for $\hat{\Pi}$ yields Equation (3).

will reflect (1) the expected future after-tax dividend stream and (2) the after-tax cash flows generated by the optimal tax-trading strategy. Given an equilibrium price, however, the objective of the investor is to follow the tax-trading strategy that maximizes his personal valuation, $W(P_t, B, b)$.

2. The Investor's Optimization Problem

In this section, we formulate the investor's problem of determining the tax-trading strategy that maximizes the value of his personal holdings of the taxable security. The basic development is similar to that found in Williams (1985). We first derive the difference equation and associated boundary conditions that describe the investor's personal valuation of the taxable security. We then show that the equilibrium price and investor's personal valuation are linear homogeneous functions and use these properties to simplify the investor's optimization problem.

In an economy with constant state prices, the equilibrium price of the taxable security will be a function of its current dividend, X_t , but will not depend explicitly on time.¹⁰ That is,

$$P_t = P(X_t), \quad (4)$$

where $P(\cdot)$ is a time-independent pricing function. This allows us to rewrite the value of the investor's position in one share of stock with current price P , basis B , and age b as follows:

$$W(P, B, b) = W(P(X), B, b) \equiv V(X, B, b), \quad (5)$$

where we have dropped the time index on X and P for convenience. In a competitive capital market, the value of a position in one share of stock at the initial purchase date for an investor who follows the optimal tax-trading strategy must equal the cost of establishing the position. Therefore, we require

$$V(X, (1+c)P(X), 0) = (1+c)P(X), \quad (6)$$

where $(1+c)P$ is the initial cost of purchasing the security and c is the

¹⁰ The equilibrium price of the taxable security is also independent of the basis values and acquisition times of those investors who currently hold the asset. This result follows from the fact that the reservation purchase price for the asset is the same across all investors. Thus, despite the possibility that investors with different basis values may have different reservation selling prices, competition among the large number of potential buyers will drive the market price of the asset equal to the common reservation purchase price. The distribution of basis values across investors will affect only the equilibrium quantity of trade. For further discussion of this point, see Dammon and Spatt (1995).

one-way transaction cost rate. We refer to Equation (6) as the “market-clearing condition” and require that it be satisfied for all values of X .

At each trading date beyond the initial purchase date, the investor has the option to either sell the stock to realize a capital gain or loss, or hold the stock to the next trading date. Suppose the investor has held the stock for b periods without a sale and assume that the investor’s basis is B , the current dividend is X , and the current price is P . If the investor optimally sells the stock at this time, he will pay a transaction cost of cP and realize a gain or loss of $(1 - c)P - B$. Thus, the value of the investor’s position in the stock at all optimal selling times is given by

$$V(X, B, b) = (1 - c)(1 - \tau)P(X) + \tau B, \quad (7)$$

where $\tau = \tau_L$ if the gain or loss is long term (i.e., if $b > N$) and $\tau = \tau_S$ if the gain or loss is short term (i.e., if $b \leq N$). At all optimal nonselling times, the investor’s personal valuation is equal to the discounted value of the one-period-ahead cum dividend payoffs on his position:

$$\begin{aligned} V(X, B, b) = & \pi_u[uX(1 - \tau_D) + V(uX, B, b + 1)] \\ & + \pi_d[u^{-1}X(1 - \tau_D) + V(u^{-1}X, B, b + 1)], \end{aligned} \quad (8)$$

where the first term in Equation (8) represents the discounted cum dividend value of the investor’s position in the stock in the “up state” and the second term represents the associated valuation in the “down state.” For each (X, B, b) combination, the investor compares the liquidation value in Equation (7) to the continuation value in Equation (8) and chooses the policy that provides the highest valuation. The equilibrium price of the taxable security is then determined by requiring that the valuation equation V satisfy the market-clearing condition given by Equation (6). This procedure ensures that the valuation function V and equilibrium pricing function P reflect the optimal realization strategy.

The solution to the investor’s optimization problem can be simplified by first identifying the form of the equilibrium pricing function P . Given the dividend process in Equation (1), the absence of arbitrage requires that the equilibrium price of the taxable security be a linear homogeneous function of its current dividend. That is,

$$P(X) = \Pi X, \quad (9)$$

where Π is the time-independent equilibrium pricing operator for the taxable security. Equation (9) implies that α shares of a stock paying a dividend of X has the same price as one share of a similar stock (with the same dividend process) paying a dividend of αX . If Equation (9) did not hold, investors could construct portfolios with positive after-

tax dividends that have a zero cost. Since this type of arbitrage is not allowed in equilibrium, Equation (9) must hold.

The linear homogeneity of the pricing function P guarantees that the investor's valuation function V is linear homogeneous in X and B . To see this, note that investors are indifferent between having α shares of stock with a current dividend of X , basis B , and age b or one share of the same stock with a current dividend of αX , basis αB , and age b . These two portfolios are equivalent because they provide exactly the same after-tax payoffs under identical tax-trading strategies. Consequently, investors' personal valuations of these two portfolios will be identical; that is, $V(\alpha X, \alpha B, b) = \alpha V(X, B, b)$ for all $\alpha > 0$.

The linear homogeneity of P and V can now be used to simplify the investor's optimization problem. Let $x \equiv X/B$ represent the dividend per dollar of basis and $v(x, b) \equiv V(x, 1, b) = V(X, B, b)/B$ represent the investor's personal valuation per dollar of basis. Then, the market-clearing condition [Equation (6)] becomes

$$v(X/(1+c)P(X), 0) = 1. \quad (10)$$

Substituting Equation (9) into Equation (10) then allows us to write the market-clearing condition as follows:

$$v(1/(1+c)\Pi, 0) = 1. \quad (11)$$

Investors' personal valuations of the taxable security for all trading dates beyond the initial purchase date can also be simplified by normalizing by the basis, B . Using Equation (8), the personal valuation per dollar of basis at all optimal nonselling times must satisfy the following linear second-order difference equation:

$$\begin{aligned} v(x, b) = & \pi_u[ux(1 - \tau_D) + v(ux, b + 1)] \\ & + \pi_d[u^{-1}x(1 - \tau_D) + v(u^{-1}x, b + 1)]. \end{aligned} \quad (12)$$

Similarly, using Equation (7), the personal valuation per dollar of basis at all optimal selling times becomes

$$v(x, b) = (1 - c)(1 - \tau)\Pi x + \tau, \quad (13)$$

where $\tau = \tau_L$ if the gain or loss is long term (i.e., if $b > N$) and $\tau = \tau_S$ if the gain or loss is short term (i.e., if $b \leq N$). The solution to Equation (12) subject to the boundary conditions [Equations (11) and (13)] is derived in the next section.

3. The Optimal Realization Policy and Equilibrium Pricing

3.1 The long-term region

Consider a position in the taxable security that is currently long term (i.e., $b > N$). Since investors have infinite horizons, the value of the position is independent of the holding period, b , once $b > N$. For example, with a one-year holding period required to qualify for long-term treatment, it does not matter whether the investor's holding period is 1 year and 1 day, or 100 years, since the investor's future after-tax payoffs are the same under identical tax-trading strategies. The time-independence of the valuation equation in the long-term region allows us to derive an analytical solution for $v(x, b)$ for all $b > N$. The solution to Equation (12) subject to the boundary condition in Equation (13) for all $b > N$ is derived in the Appendix. The solution is¹¹

$$v(x) = \{[(1-c)(1-\tau_L)\Pi - (1-\tau_D)\hat{\Pi}]x_L + \tau_L\}(x/x_L)^m + (1-\tau_D)\hat{\Pi}x \quad (14)$$

for all $b > N$ and $x > x_L$, where m is a constant given by

$$m = \{\ln[1 - (1 - 4\pi_u\pi_d)^{1/2}] - \ln[\pi_u] - \ln[2]\} / \ln[u] < 0 \quad (15)$$

and x_L defines the investor's long-term realization boundary such that for all $x \leq x_L$ the investor realizes the associated long-term gain or loss and for all $x > x_L$ the investor defers the associated long-term gain or loss.

The interpretation of Equation (14) is straightforward. The last term in Equation (14), $(1-\tau_D)\hat{\Pi}x$, is the value of the after-tax dividend income on the taxable security, where $\hat{\Pi}x$ is the value of the tax-exempt counterpart security with a current dividend of x . If all realized capital gains and losses were tax-exempt, then $(1-\tau_D)\hat{\Pi}x$ would be the investor's personal valuation of his current holdings of the taxable security. The first term in Equation (14) is the value of the investor's tax-timing option, which depends on the current dividend per dollar of basis, x , in relation to the investor's chosen long-term realization boundary, x_L .¹² The investor chooses x_L to maximize his personal valuation, $v(x)$, taking as given the market price of the taxable security.

¹¹ Since $v(x, b)$ is independent of b for all $b > N$, we have written v as a function only of x in Equation (14).

¹² The term within the set brackets in Equation (14) is the value of the investor's tax-timing option when $x = x_L$. It is composed of the after-tax proceeds from an immediate sale, $(1-c)(1-\tau_L)\Pi x_L + \tau_L$, less the value of the after-tax dividend stream, $(1-\tau_D)\hat{\Pi}x_L$. Multiplying this tax-timing option value by $(x/x_L)^m$ gives the value of the investor's tax-timing option for values of $x > x_L$.

Of course, the optimal realization policy will ultimately influence the price of the taxable security through the market-clearing condition, but as a price-taker the investor ignores the impact of his actions on the equilibrium price.

Equation (14) implies that if $[(1 - c)(1 - \tau_L)\Pi - (1 - \tau_D)\hat{\Pi}] > 0$, then $v(x)$ is globally increasing in x_L . In this case, the optimal tax-trading policy is to realize all gains and losses in the long-term region (i.e., $x_L = \infty$) and the value of the investor's position in the long-term region is given by Equation (13) (with $\tau = \tau_L$) for all values of x . If $\tau_L > 0$ and $[(1 - c)(1 - \tau_L)\Pi - (1 - \tau_D)\hat{\Pi}] < 0$, the optimal tax-trading policy will involve the deferral of some long-term gains and, depending upon the level of transaction costs, may also involve the deferral of some long-term losses (i.e., $0 < x_L < \infty$). If $\tau_L = 0$ and $[(1 - c)(1 - \tau_L)\Pi - (1 - \tau_D)\hat{\Pi}] < 0$, which requires a sufficiently high level of transaction costs, the right-hand side of Equation (14) is globally decreasing in x_L . Hence, the optimal tax-trading policy entails the deferral of *all* long-term gains and losses (i.e., $x_L = 0$) and the value of the investor's position in the long-term region reduces to $v(x) = (1 - \tau_D)\hat{\Pi}x$. In this case, the transaction costs of trading the taxable security exceed the tax benefits of resetting the tax basis and restarting the short-term holding period.

Since x follows a discrete binomial process, any interior solution for x_L must lie on the discrete grid that describes feasible values for x . Nevertheless, useful insights about the optimal solution can be obtained by ignoring, for the moment, the restriction that x_L lie on the discrete grid. If the optimal realization boundary in the long-term region could take on any value in the interval $[0, \infty)$, then it would be characterized by setting the derivative of Equation (14) with respect to x_L equal to zero. This yields

$$x_L = \frac{\tau_L m}{[(1 - c)(1 - \tau_L)\Pi - (1 - \tau_D)\hat{\Pi}](1 - m)}, \quad (16)$$

provided $\tau_L > 0$ and $[(1 - c)(1 - \tau_L)\Pi - (1 - \tau_D)\hat{\Pi}] < 0$.¹³ The

¹³ Substituting Equation (16) into Equation (14) and simplifying yields the following expression for the investor's personal valuation in the long-term region for all $x > x_L$:

$$v(x) = \left(\frac{\tau_L}{1 - m} \right) \left(\frac{x}{x_L} \right)^m + (1 - \tau_D)\hat{\Pi}x,$$

where m is given by Equation (15) and x_L is given by Equation (16). In the special case of symmetric taxation (i.e., $\tau_L = \tau_S$) without transaction costs, the optimal long-term cutoff level applies at all dates and is given by $x_L = 1/\Pi$; that is, investors optimally realize all losses and defer all gains. Letting $x = x_L$ and realizing that $v(1/\Pi) = 1$ by the market-clearing condition then yields:

$$1 = \tau_L/(1 - m) + (1 - \tau_D)\hat{\Pi}/\Pi.$$

second derivative of Equation (14) evaluated at x_L is strictly negative, ensuring that the solution described by Equation (16) is a maximum on the continuous grid. The optimal long-term realization boundary on the discrete grid can be found by considering the values of x that straddle the continuous solution given by Equation (16) and choosing the one that produces the highest value for the taxable security. The details of this procedure are described later in this section.

The optimal long-term realization boundary also can be expressed as the ratio of the price to the basis, P/B , above which it is optimal to defer all long-term capital gains and losses. This optimal long-term realization boundary can be written as follows:

$$P/B = \Pi x_L = \frac{\tau_L m}{[(1-c)(1-\tau_L) - (1+c)(1-z)](1-m)}, \quad (17)$$

where

$$z \equiv 1 - \frac{\hat{\Pi}(1-\tau_D)}{\Pi(1+c)} \quad (18)$$

is the value of the tax-timing option per dollar invested in the taxable security.¹⁴ Since the tax payments and rebates on optimal realizations are proportional to the original investment in the taxable security, so too is the value of the tax-timing option.

3.2 The short-term region

Having derived the optimal realization policy and personal valuation equation for the long-term region, we now turn our attention to the short-term region. Because the investor's personal valuation of the taxable security is time dependent in the short-term region, we cannot derive the optimal solution analytically. Instead, we numerically derive the optimal realization policy and the investor's personal valuation in the short-term region using dynamic programming. The market-

The first term on the right-hand side represents the value of the tax-timing option per dollar invested in the taxable security and is identical to the expression derived by Constantinides [1983, Equation (21)]. The second term on the right-hand side represents the value of the after-tax dividends per dollar invested in the taxable security. Multiplying through by Π and rearranging terms gives an expression for the equilibrium pricing operator for the taxable security under symmetric taxation (without transaction costs):

$$\Pi = \frac{(1-\tau_D)\hat{\Pi}}{1-\tau_L/(1-m)}.$$

The equilibrium pricing operator for the taxable security is higher (lower) than that for the tax-exempt security if τ_D is less than (greater than) $\tau_L/(1-m)$.

¹⁴ The cost of acquiring the taxable security, including transaction costs, is $\Pi X(1+c)$. This cost is equal to the sum of the value of the after-tax dividends, $\hat{\Pi}X(1-\tau_D)$, and the value of the tax-timing option, $z\Pi X(1+c)$. Solving this relationship for z gives the desired expression.

clearing condition is then used to identify the equilibrium price of the taxable security.

The equilibrium pricing operator for the taxable security, Π , and the optimal realization policy are determined through an iterative procedure. We first choose an initial value of Π as a candidate for the equilibrium pricing operator. Given this value of Π , we then determine the optimal realization policy in the long-term region (i.e., for $b > N$) using Equation (16). The investor's personal valuation of the taxable security at the first date in the long-term region (i.e., for $b = N + 1$) is then determined for all $x \leq x_L$ using Equation (13) (with $\tau = \tau_L$) and for all $x > x_L$ using Equation (14). These valuations are then used to solve recursively for the valuations and optimal realization policy in the short-term region using Equations (12) and (13) (with $\tau = \tau_S$). At each date and for each state in the short-term region, a comparison is made between the values of realizing and deferring the associated gain or loss. The policy that provides the highest value for $v(x, b)$ defines the optimal strategy at each date and for each state. We then check to see whether the investor's personal valuation at the initial purchase date (at $b = 0$) satisfies the market-clearing condition given by Equation (11). If not, a new value of Π is chosen by linear interpolation and the valuation problem is solved again. This procedure continues until the value of Π that satisfies the market-clearing condition is found.

The above procedure is complicated by the fact that the optimal realization cutoff level in the long-term region, x_L , must lie on the discrete grid of feasible values for x . Since the optimal realization boundary defined by Equation (16) is not restricted to lie on the discrete grid, it cannot be used directly in the valuation of the investor's position in the taxable security in the long-term region. However, since the valuation function is uniquely maximized at the continuous value of x_L defined by Equation (16), the optimal realization boundary on the discrete grid must either be at the node immediately below or at one of the two nodes immediately above this continuous value.¹⁵ Thus, for each of the three possible optimal values of x_L on the discrete grid, we carry out the procedure outlined above to determine the equilibrium pricing operator for the taxable security, Π . The optimal realization boundary on the discrete grid is then defined to be that value of x_L that produces the highest value of Π .

¹⁵ The optimal cutoff level on the discrete grid is defined to be the highest value of x for which it is optimal to realize the position. Because of the nature of the continuous solution, we know that the optimal cutoff on the discrete grid must be at, or above, the node immediately below the continuous cutoff value. The need to consider the two nodes immediately above the continuous cutoff value stems from the fact that the discrete grid differs at even and odd dates.

4. Numerical Examples

4.1 Optimal long-term realization policy and the value of tax timing

In this section we use some numerical examples to illustrate the optimal trading and equilibrium pricing of the taxable security. Consider an economy with a taxable and tax-exempt security that both pay dividends according to the same binomial process described by Equation (1). The mean and variance of this process are

$$\mu - 1 = E[X_{t+1}/X_t] - 1 = qu + (1 - q)u^{-1} - 1 \quad (19)$$

and

$$\sigma^2 = \text{var}[X_{t+1}/X_t] = (u - \mu)(\mu - u^{-1}), \quad (20)$$

where q is the probability of the “up state” and $(1 - q)$ is the probability of the “down state.” Solving Equation (20) for u , recognizing that $u > 1$, yields

$$u = \frac{(1 + \mu^2 + \sigma^2) + [(1 + \mu^2 + \sigma^2)^2 - 4\mu^2]^{1/2}}{2\mu}. \quad (21)$$

Solving Equation (19) for the probability of the “up state,” q , yields

$$q = (\mu - u^{-1})/(u - u^{-1}). \quad (22)$$

Substituting the above value for q into the equilibrium state prices for a risk-neutral economy, $\pi_u = qR^{-1}$ and $\pi_d = (1 - q)R^{-1}$, yields

$$\pi_u = (\mu - u^{-1})/R(u - u^{-1}) \quad (23)$$

and

$$\pi_d = (u - \mu)/R(u - u^{-1}), \quad (24)$$

where R is one plus the constant tax-exempt, riskless interest rate. Given the parameters u , π_u , and π_d , the equilibrium pricing operator for the tax-exempt security, $\hat{\Pi}$, is given by Equation (3).

We initially assume that dividends and price changes occur weekly, with a mean and standard deviation of $\mu - 1 = 0.09387\%$ per week (5% per annum) and $\sigma = 4.5\%$ per week (32.45% per annum), respectively. Since the equilibrium prices of the taxable and tax-exempt securities are a constant proportion of the dividend, μ and σ are also the mean and standard deviation of the weekly capital gain return on the taxable and tax-exempt securities. Our choice of $\sigma = 4.5\%$ per week is an attempt to approximate the standard deviation of the capital gain return for an average stock. The tax-exempt, riskless interest rate is assumed to be $R - 1 = 0.18346\%$ per week (10% per annum).

Using these parameter values, the equilibrium pricing operator for the tax-exempt security is $\hat{\Pi} = 1117.3586$.

Table 1 provides the results for a range of values for c (the one-way transaction cost rate), τ_L (the long-term tax rate), and N (the length of the short-term region). These results are based on the assumption that the short-term tax rate is $\tau_s = 40\%$ and the dividend tax rate is $\tau_D = 0\%$. A dividend tax rate of zero allows us to focus exclusively on the effects of capital gains taxation and transaction costs on the relative pricing of the taxable and tax-exempt securities. Table 1 reports the ratio of the price of the taxable security to the price of its tax-exempt counterpart, $\Pi/\hat{\Pi}$, the value of the tax-timing option per dollar invested in the taxable security, z , and the optimal long-term realization boundary in terms of the ratio of the price to the basis, P/B , above which it is optimal to defer the long-term capital gain or loss.

Table 1 indicates that transaction costs can dramatically affect the optimal long-term realization boundary and the value of the tax-timing option. For example, with $N = 52$ weeks and $\tau_L = 28\%$, the investor optimally realizes all long-term capital gains below 222.1% (i.e., $P/B = 3.221$) in the absence of transaction costs, but realizes only those capital gains below 4.1% (i.e., $P/B = 1.041$) with $c = 0.5\%$. The value of the tax-timing option declines by over 46% in this case. Thus, for a relatively small increase in transaction costs, the taxable security's equilibrium market price and optimal long-term realization boundary decline dramatically. Subsequent increases in the level of transaction costs have a much smaller marginal impact on the value of the tax-timing option and the optimal long-term realization boundary.¹⁶

The effect of an increase in the long-term tax rate, τ_L , on the optimal trading and equilibrium pricing of the taxable security depends upon the level of transaction costs. For sufficiently low transaction costs, investors optimally realize some long-term capital gains to restart the short-term holding period (i.e., $P/B > 1$). In this case, an increase in the long-term tax rate increases the cost of realizing long-term capital gains and reduces the value of the tax-timing option. For sufficiently

¹⁶ The dramatic impact of transaction costs on the optimal realization behavior is further illustrated by the case of symmetric taxation. From the last three columns of Table 1, it can be seen that investors optimally realize all losses and defer all gains (i.e., $P/B = 1.0$) in the absence of transaction costs, but optimally realize only those losses greater than 13.1% (i.e., $P/B = 0.869$) with $c = 0.5\%$. Thus, with $c = 0.5\%$, investors defer the realization of capital losses until the tax rebate on the loss is about six times as large as the pretax round-trip transaction costs. This realization behavior reflects the dynamic nature of the realization problem in which investors attempt to maximize the present value of the difference between the tax rebates and transaction costs over time. A similar trade-off between interest savings and refinancing costs arises in the mortgage refinancing problem [see Dunn and Spatt (1986)]. The above intuition also carries over to the case with asymmetric capital gains taxes.

Table 1
Value of the tax-timing option and the optimal long-term realization boundary for medium variance stocks^a

c^b	$\tau_L = 20\%^c$			$\tau_L = 28\%$			$\tau_L = 40\%$		
	$\Pi/\hat{\Pi}^d$	z^e	P/B^f	$\Pi/\hat{\Pi}$	z	P/B	$\Pi/\hat{\Pi}$	z	P/B
Panel A: $N = 26$ weeks ^g									
0.0%	6.818	0.853	∞	1.619	0.382	∞	1.202	0.168	1.000
0.5%	1.416	0.297	∞	1.144	0.130	0.995	1.159	0.142	0.869
1.0%	1.102	0.102	0.990	1.112	0.110	0.865	1.139	0.131	0.791
2.0%	1.058	0.074	0.716	1.077	0.089	0.716	1.110	0.116	0.716
Panel B: $N = 52$ weeks									
0.0%	2.306	0.566	∞	1.295	0.228	3.221	1.202	0.168	1.000
0.5%	1.306	0.238	∞	1.151	0.136	1.041	1.159	0.142	0.869
1.0%	1.121	0.116	1.133	1.122	0.118	0.905	1.139	0.131	0.791
2.0%	1.074	0.087	0.783	1.086	0.098	0.748	1.110	0.116	0.716
Panel C: $N = 104$ weeks									
0.0%	1.582	0.368	∞	1.220	0.180	1.568	1.202	0.168	1.000
0.5%	1.202	0.172	3.205	1.155	0.139	1.089	1.159	0.142	0.869
1.0%	1.132	0.125	1.240	1.130	0.124	0.905	1.139	0.131	0.791
2.0%	1.089	0.099	0.857	1.096	0.105	0.783	1.110	0.116	0.716

^aThe values in the table are based on the assumptions that investors are risk neutral, the mean growth rate in dividends is $\mu - 1 = .09387\%$ per week (about 5% per annum), the standard deviation of the growth rate in dividends and capital gains is $\sigma = 4.5\%$ per week (about 32.45% per annum), the tax-exempt, riskless interest rate is $R - 1 = 0.18346\%$ per week (about 10% per annum), the short-term tax rate is $\tau_s = 40\%$, and the dividend tax rate is $\tau_D = 0\%$.

^b c is the one-way transaction cost rate.

^c τ_L is the long-term tax rate.

^d $\Pi/\hat{\Pi}$ is the ratio of the price of the taxable security to the price of its tax-exempt counterpart.

^e z is the value of the tax-timing option per dollar invested in the taxable security.

^f P/B is the optimal long-term cutoff level describing the ratio of the price to the basis above which all long-term capital gains and losses are deferred and below which all long-term capital gains and losses are realized.

^g N is the number of trading periods in the short-term region.

high transaction costs, investors optimally defer all long-term capital gains and only realize larger long-term capital losses (i.e., $P/B < 1$). In this case, an increase in the long-term tax rate benefits the investor by increasing the tax rebates on his long-term capital loss realizations. This increases the value of the tax-timing option.¹⁷

¹⁷ Table 1 indicates that as the long-term tax rate is reduced, the optimal long-term cutoff level, P/B , increases. This behavior can be understood as follows. When $P/B > 1$, a lower long-term tax rate reduces the cost of realizing long-term capital gains, thereby giving investors the incentive to realize larger long-term capital gains. When $P/B < 1$, a lower long-term tax rate reduces the tax rebates on long-term capital loss realizations, thus making short-term status relatively more valuable. As a consequence, investors are relatively more aggressive at realizing long-term capital losses to restart the short-term holding period. Moreover, investors are also more aggressive at realizing short-term capital losses to avoid long-term tax treatment.

With symmetric taxation, the length of the short-term region, N , is irrelevant for the optimal trading and equilibrium pricing of the taxable security. With asymmetric taxation, the effect of an increase in the length of the short-term region on the optimal trading and equilibrium pricing of the taxable security depends upon the level of transaction costs. In the absence of transaction costs, a shorter short-term region is unambiguously more valuable since it allows the investor to exploit the asymmetry between the long-term and short-term tax rates more frequently. Consequently, the value of the tax-timing option and the optimal long-term realization boundary are inversely related to N in the absence of transaction costs. With positive transaction costs, however, a shorter short-term region may not be more valuable to the investor. For example, in the extreme case where the short-term region is a single day, the transaction costs of realizing short-term losses will typically exceed the tax rebate on the loss. In this case, there is no incentive for the investor to realize long-term gains to restart the short-term holding period. At the other extreme, however, a short-term region that lasts for 100 years offers little opportunity for the investor to exploit the asymmetry between the long-term and short-term tax rates. Consequently, in the presence of transaction costs, the value of the tax-timing option and the optimal long-term realization boundary will typically be maximized at some intermediate value of N .

Although not shown in Table 1, the value of the tax-timing option and the optimal long-term realization boundary are unambiguously increasing in the level of the short-term tax rate, τ_s . This reflects the fact that a higher short-term tax rate increases the tax rebates on the realization of short-term capital losses. An increase in the volatility of the stock (as shown below in Table 3 under the heading "Optimal Tax Trading") unambiguously increases the value of the tax-timing option. This behavior is consistent with standard option pricing results. In the absence of transaction costs, an increase in the volatility also increases the optimal long-term realization boundary, P/B . However, with sufficiently high transaction costs, the optimal long-term realization boundary may be less than one and, in this case, may not be strictly increasing in the volatility.¹⁸ Table 2 summarizes our comparative static results.

¹⁸ Recall that with transaction costs, investors attempt to maximize the present value of the difference between the tax rebates and transaction costs over time. Thus, investors are reluctant to realize small current losses for fear of losing the opportunity to realize substantially larger losses in the future without incurring additional transaction costs. Since large future losses are more likely to occur with higher volatility, the optimal long-term realization boundary can be lower for higher volatility stocks in the presence of transaction costs. This phenomenon occurs in Table 3 under the Optimal Tax Trading strategy for the medium-variance and low-variance stocks with one-way transaction costs of 2%.

Table 2
Summary of comparative statics

Parameter ^a	Impact on:		
	$\Pi/\hat{\Pi}^b$	z^c	P/B^d
c	—	—	—
σ	+	+	$\pm$
N	$\pm$	$\pm$	$\pm$
τ_D	—	0	0
τ_L	$\pm$	$\pm$	$\pm$
τ_S	+	+	+

^a c is the one-way transaction cost rate, σ is the standard deviation of the growth rate in dividends and capital gain return, N is the length of the short-term region, τ_D is the dividend tax rate, τ_L is the long-term tax rate, and τ_S is the short-term tax rate.

^b $\Pi/\hat{\Pi}$ is the ratio of the price of the taxable security to the price of its tax-exempt counterpart.

^c z is the value of the tax-timing option per dollar invested in the taxable security.

^d P/B is the optimal long-term cutoff level describing the ratio of the price to the basis above which all long-term capital gains and losses are deferred and below which all long-term capital gains and losses are realized.

The overall impression given by Table 1 is that the value of the tax-timing option can be considerable. For example, with $\tau_L = 28\%$ and $N = 52$ weeks, the value of the tax-timing option represents nearly 23% of the total market value for medium-variance stocks in the absence of transaction costs. With transaction costs of $c = 1\%$, the value of the tax-timing option for medium-variance stocks still accounts for nearly 12% of the total market value. The corresponding values for high-variance stocks ($\sigma = 9\%$ per week) are 69.7% with $c = 0\%$ and 24.2% with $c = 1\%$ (see Table 3). The value of the tax-timing option for low-volatility stocks is typically small for all levels of transaction costs, except in cases where the long-term tax rate is considerably less than the short-term tax rate.

4.2 Optimal trading in the short-term region

The behavior of the optimal realization boundary in the short-term region is also of considerable interest. To illustrate how the optimal realization boundary in the short-term region behaves, we examine a special case in which dividends, stock price changes, and trading occur daily. Although the assumption of daily trading is made so that the effects of discrete price changes on the optimal short-term realization boundary are minimal, the qualitative features of the short-term realization boundary would also arise under alternative assumptions about the frequency of dividends, stock price changes and trading. We assume that the length of the short-term region is $N = 260$ days (i.e., one trading year), the short-term tax rate is $\tau_S = 40\%$, the long-term tax rate is $\tau_L = 28\%$, the dividend tax rate is $\tau_D = 0\%$, the one-way transaction cost rate is $c = 0\%$, the mean growth rate in dividends is




Figure 1
Optimal short-term cutoff level

$\mu - 1 = 0.01877\%$ per day (5% per annum), the standard deviation of the growth rate in dividends is $\sigma = 2\%$ per day (32.25% per annum), and the tax-exempt, riskless interest rate is $R - 1 = 0.0367\%$ per day (10% per annum). Using these parameter values, the equilibrium pricing operator for the tax-exempt security is $\hat{\Pi} = 5592.8145$.

Under these conditions, the equilibrium value of the taxable security relative to its tax-exempt counterpart and the value of the tax-timing option per dollar invested in the taxable security are $\Pi/\hat{\Pi} = 1.328$ and $z = 0.247$, respectively. The optimal realization policy in the long-term region is to realize all capital gains below 395.2% (i.e., $P/B = 4.952$ for all $b > 260$). The optimal realization boundary at each date in the short-term region is shown in Figure 1. All short-term capital gains and losses above the realization boundary depicted in Figure 1 are optimally deferred, while those at or below the realization boundary are optimally realized.

As Figure 1 indicates, the optimal realization boundary in the short-term region is time dependent. The sawtooth pattern exhibited by the optimal short-term realization boundary reflects the fact that the dividend and equilibrium price follow a discrete binomial process. With the exception of the last date in the short-term region (i.e., day 260), the optimal short-term realization boundary lies strictly *below* the point at which the price and basis are equal. The maximum loss that is optimally deferred is a function of the time remaining in the short-term region. For the example illustrated in Figure 1, the opti-

mal trading policy involves the deferral of short-term losses as large as 10%, even though transaction costs are zero. For higher volatility stocks, the optimal realization policy involves the deferral of even larger short-term losses. These results are striking and are in direct opposition to the common intuition that investors should optimally realize all short-term capital losses as soon as they occur in the absence of transaction costs.

The optimality of deferring small short-term losses stems from the fact that restarting the short-term holding period prior to reaching the long-term boundary can be disadvantageous for the investor. For example, consider an investor with a short-term position in a stock whose current price equals his basis. If the investor sells and repurchases the stock, there is no immediate tax liability for the investor, but the short-term holding period is restarted. This is disadvantageous for the investor since it extends the length of time he must wait before receiving the more favorable long-term tax treatment on any *subsequent* capital gains. In fact, the disadvantage of restarting the short-term holding period prior to reaching the long-term region can exceed the tax benefits of realizing small short-term capital losses. Of course, at the final date in the short-term region there remains no benefit to deferring the realization of a capital loss that is about to become long-term. Hence, the optimal realization boundary at the final date in the short-term region is $P/B = 1$.¹⁹ These results indicate that under the optimal tax-trading policy investors may be quick to realize long-term gains to restart the short-term holding period, but may be slow to realize short-term losses.

4.3 Extensions to multiple holding periods

Prior to 1942 in the U.S., and at various times in other countries, the tax rate on capital gains and losses varied over three or more separate holding periods. The dynamic programming approach that is used to determine the equilibrium price of the taxable security and the optimal realization policy can also be used to evaluate these more complicated tax environments. For example, consider the case where the capital

¹⁹ In the presence of transaction costs, the optimal realization boundary in the short-term region will still fall below $P/B = 1$, but may not exhibit the U-shape pattern illustrated in Figure 1. The reason is that if transaction costs are high enough so that investors optimally defer the realization of all long-term capital gains, then restarting the short-term holding period by realizing short-term capital losses can be advantageous for the investor. In this case, the optimal realization boundary in the short-term region will be influenced primarily by investors' desire to avoid transaction costs. Under the optimal trading policy, the amount by which the tax rebate on loss realizations must exceed transaction costs is an increasing function of the time remaining in the short-term region. In other words, to avoid paying excessive transaction costs, investors will optimally defer larger losses the more time that remains in the short-term region.

gains tax rate varies over three separate holding periods:

$$\tau = \begin{cases} \tau_S, & \text{if } b \leq N_1 \\ \tau_M, & \text{if } N_1 < b \leq N_2 \\ \tau_L, & \text{if } b > N_2, \end{cases} \quad (25)$$

where $\tau_S > \tau_M > \tau_L$ and $N_1 < N_2$. Since our analysis of the long-term region (i.e., for $b > N_2$) is independent of the number of separate holding periods, the optimal long-term realization boundary is again given by Equation (16) and the value of the investor's position in the taxable security for each node in the long-term region can be determined by Equations (13) and (14). The valuations and optimal realization policies in the short-term and intermediate-term regions can then be determined by using the dynamic programming approach described earlier, with the capital gains tax rate for these two regions given by Equation (25).

To illustrate, let $\tau_S = 40\%$, $\tau_M = 28\%$, $\tau_L = 15\%$, $N_1 = 52$ weeks (i.e., a 1-year short-term region) and $N_2 = 156$ weeks (i.e., a 2-year intermediate-term region). Assets held longer than 156 weeks qualify for long-term tax treatment. We assume that the stock generates dividends and price changes weekly. The parameters of the dividend and stock price process are assumed to be the same as those used to construct Table 1. Transaction costs and the dividend tax rate are assumed to be zero. Under these conditions $\Pi/\hat{\Pi} = 1.664$, $z = 0.399$, and $P/B = \infty$ for all $b > 156$ weeks.

The optimal realization boundary in the short-term and intermediate-term regions is shown in Figure 2. The behavior of the optimal realization boundary in the short-term region (i.e., for $b \leq 52$ weeks) is similar to the behavior exhibited in Figure 1 and reflects the reluctance to restart the short-term holding period prematurely. The realization of a capital gain or loss in the intermediate-term region has both a positive and a negative effect on the investor. On the positive side, the realization of an intermediate-term capital gain or loss allows the investor to reset his tax basis and restart the option to realize subsequent capital losses at the higher short-term tax rate, τ_S . On the negative side, restarting the short-term holding period increases the length of time the investor must wait to realize subsequent capital gains at the more favorable long-term tax rate, τ_L . At the start of the intermediate-term region, the first effect dominates and the investor is willing to realize some small capital gains to reestablish the short-term holding period. As the investor moves further into the intermediate-term region, however, the second effect becomes increasingly more important and the optimal intermediate-term realization boundary declines. In fact, Figure 2 indicates that it can be optimal to defer the realization of small




Figure 2
Optimal short-term and intermediate-term cutoff levels

capital losses over some portion of the intermediate-term region to avoid restarting the short-term holding period. At the last date in the intermediate region, the optimal realization boundary again reaches $P/B = 1$.

5. Comparisons With Alternative Tax-Trading Strategies

As the previous examples illustrate, the optimal tax-trading policy is complicated in the presence of asymmetric taxation and transaction costs. This raises the question as to whether the optimal tax-trading policy provides significantly higher benefits than some simpler tax-trading strategies that are not optimal, but may be easier to implement in practice. To investigate this question, we examine two alternative tax-trading strategies: (1) simplified tax trading and (2) restricted tax trading. Under both strategies, trading is limited to at most once per year, which reduces the time and effort investors must spend in monitoring portfolio positions.

Simplified tax trading assumes that investors (1) defer all short-term capital gains, (2) realize all short-term losses (and break-even positions) at the last date in the short-term region, but do not trade prior

to that time, (3) realize all long-term capital losses (and break-even positions) as soon as they occur, and (4) either realize all long-term capital gains each year, or defer all long-term capital gains indefinitely, depending upon which provides the highest value for the taxable security. Simplified tax trading is intended to approximate, within our framework, the realization strategy proposed by Constantinides (1984).

Restricted tax trading assumes that investors trade the taxable security optimally in both the long-term and short-term regions, subject to the restriction that investors are only allowed to trade at the last date in the short-term region, but may trade at any date in the long-term region. Restricted tax trading is identical to simplified tax trading, except that it allows the investor to choose the short-term and long-term realization boundaries optimally (subject to the trading restrictions described above).

Table 3 provides a comparison of the results for simplified tax trading, restricted tax trading, and optimal tax trading. Optimal tax trading is the trading strategy that is optimal in our environment when investors are allowed to trade at each date in both the short-term and long-term regions. The results are reported for three variance groups and for a range of transaction costs. Our results for simplified tax trading indicate that realizing long-term capital gains each year outperforms the policy of deferring long-term capital gains only for high-variance stocks with zero transaction costs. This result is consistent with the theoretical results reported by Constantinides (1984) and the empirical findings of Dammon, Dunn, and Spatt (1989).²⁰ Restricted tax trading produces a value for the tax-timing option that is at least as high as that produced by simplified tax trading, although the differences are typically small. This suggests that the incremental value of allowing investors to optimally choose the long-term realization boundary is relatively small when investors are restricted from trading prior to the last date in the short-term region.

Optimal tax trading allows investors to trade optimally at each date in both the long-term and short-term regions. Under optimal tax trading, investors are more willing to realize long-term capital gains (or losses) to restart the short-term holding period due to the fact that trading in the short-term region is unrestricted under this strategy. The ability to trade optimally at each date in the short-term region also

²⁰ Constantinides (1984, Table 1) reports that, with $\tau_L/\tau_S = 0.7$ and $R = 1.10$, the optimal realization policy in the absence of transaction costs is to realize all long-term capital gains each year provided the standard deviation of the annual stock return is at least 60%. With one-way transaction costs of 2%, Constantinides (1984, Table 2) reports that it is optimal to realize long-term capital gains each year only if the standard deviation of the annual stock return is at least 100%.

Table 3
Comparisons of alternative tax-trading policies^a

c^f	Simplified tax trading ^b			Restricted tax trading ^c			Optimal tax trading ^d			Lower bound values with offset rules ^e		
	$\Pi/\hat{\Pi}^g$	z^h	P/B^i	$\Pi/\hat{\Pi}$	z	P/B	$\Pi/\hat{\Pi}$	z	P/B	$\Pi/\hat{\Pi}$	z	P/B
Panel A: High variance ($\sigma = 9\%$ per week)												
0.0%	1.390	0.281	∞	1.390	0.281	∞	3.302	0.697	∞	1.394	0.282	1.000
0.5%	1.284	0.225	0.995	1.288	0.228	1.303	1.594	0.376	∞	1.334	0.254	0.995
1.0%	1.259	0.214	0.990	1.259	0.214	0.990	1.307	0.242	1.552	1.298	0.237	0.990
2.0%	1.212	0.191	0.980	1.221	0.197	0.684	1.244	0.212	0.819	1.244	0.212	0.819*
Panel B: Medium variance ($\sigma = 4.5\%$ per week)												
0.0%	1.146	0.128	1.000	1.147	0.128	1.046	1.295	0.228	3.221	1.202	0.168	1.000
0.5%	1.122	0.113	0.995	1.124	0.114	0.909	1.151	0.136	1.041	1.151	0.135	0.995
1.0%	1.099	0.099	0.990	1.106	0.105	0.827	1.122	0.118	0.905	1.122	0.118	0.905*
2.0%	1.055	0.070	0.980	1.078	0.091	0.716	1.086	0.098	0.748	1.086	0.098	0.748*
Panel C: Low variance ($\sigma = 2.5\%$ per week)												
0.0%	1.058	0.055	1.000	1.058	0.055	1.000	1.096	0.088	1.133	1.090	0.083	1.000
0.5%	1.039	0.042	0.995	1.043	0.046	0.900	1.055	0.056	0.946	1.055	0.056	0.946*
1.0%	1.030	0.030	0.990	1.031	0.040	0.852	1.037	0.046	0.874	1.037	0.046	0.874*
2.0%	0.985	0.005	0.980	1.011	0.031	0.783	1.014	0.033	0.783	1.014	0.033	0.783*

^aThe values in the table are based on the assumptions that investors are risk neutral, the mean growth rate in dividends is $\mu - 1 = 0.09387\%$ per week (about 5% per annum), the tax-exempt, riskless interest rate is 0.18346% per week (about 10% per annum), the length of the short-term holding period is $N = 52$ weeks, the dividend tax rate is $\tau_D = 0\%$, the short-term tax rate is $\tau_S = 40\%$, and the long-term tax rate is $\tau_L = 28\%$.

^bSimplified tax trading assumes that investors (1) defer all short-term gains, (2) realize all short-term losses at the last date in the short-term region, but do not trade prior to that time, (3) realize all long-term losses as soon as they occur, and (4) either realize all long-term capital gains each year, or defer all long-term capital gains indefinitely, depending upon which provides the highest value for the taxable security.

^cRestricted tax trading assumes that investors trade the taxable security optimally subject to the restriction that they are allowed to trade only at the last date in the short-term region, but are unrestricted in the long-term region.

^dOptimal tax trading assumes that investors trade the taxable security optimally at each trading date in both the long-term and short-term regions.

^eThe lower bound values with offset rules are the values for the optimal tax trading policy assuming that investors do not realize long-term capital gains.

^f c is the one-way transaction cost rate.

^g $\Pi/\hat{\Pi}$ is the ratio of the price of the taxable security to the price of its tax-exempt counterpart.

^h z is the value of the tax-timing option per dollar invested in the taxable security.

ⁱ P/B is the long-term cutoff level describing the ratio of the price to the basis above which all long-term capital gains and losses are deferred and below which all long-term capital gains and losses are realized.

*Indicates that offset rules have no effect on the optimal realization policy or the value of tax timing.

adds significant value to the tax-timing option (see Table 3), especially for high-variance stocks with low transaction costs. For example, the value of the tax-timing option is 489.2% higher under optimal tax trading than it is under restricted tax trading for high-variance stocks with zero transaction costs. With one-way transaction costs of $c = 1\%$, the corresponding percentage increase in the value of the tax-timing option is 17.4%.

In practice, the tax code requires investors to offset net short-term capital losses with net long-term capital gains realized in the same tax year. This feature of the tax code reduces the value of realizing long-term capital gains to restart the short-term holding period since it nullifies the asymmetry between the long-term and short-term tax rates. The results reported in Table 3 for the three trading policies discussed above ignore the offset provisions of the tax code. The effect of the offset provisions on the value of tax timing is most severe for those trading strategies that involve realizing relatively large long-term capital gains. However, for trading strategies that involve the deferral of all long-term capital gains (i.e., $P/B < 1.0$), the offset provisions are of no consequence.

For those cases in Table 3 that involve the realization of some long-term capital gains (i.e., $P/B > 1.0$), the introduction of offset provisions will reduce the value of tax timing to some degree. This is true for all three trading strategies presented in Table 3, but is perhaps most severe for the optimal tax trading strategy.²¹ The magnitude of the reduction in the value of the tax-timing option is difficult to measure precisely due to the fact that the optimal realization policy in the presence of offset rules is complicated and not well understood. However, we can derive a lower bound on the value of tax timing in the presence of offset rules by assuming that investors never realize long-term capital gains. We do this for the optimal tax trading strategy by imposing the long-term realization boundary of $P/B = 1/(1 + c)$ for those cases in Table 3 for which $P/B > 1.0$ in the absence of offset rules. This allows investors to trade at each date in both the long-term and short-term regions, subject to the restriction that they cannot realize long-term capital gains. The results are reported in the last three columns of Table 3.

²¹ The optimal tax trading strategy allows investors to trade at each date in both the short-term and long-term regions. This can lead to the realization of short-term losses and long-term gains in the same tax year, which would then be subject to the offset provisions of the tax code. Under simplified tax trading, investors realize losses only at the last date in the short-term region, and either realize or defer all long-term capital gains each year. While this avoids the realization of gains and losses in the same tax year for an individual stock, it does not overcome the offset problem with a portfolio of stocks. In a portfolio context, an investor who follows the simplified tax trading strategy may realize short-term capital losses on some stocks and long-term capital gains on other stocks in the same tax year. This is also true for the restricted tax trading strategy.

Despite the fact that we are placing no value on the restarting option, the lower bounds for the value of tax timing in the presence of offset rules are still quite high. For high-variance stocks, the ability to take losses increases the value of the taxable security (over its tax-exempt counterpart) by 39.4% in the absence of transaction costs and by 29.8% in the presence of 1% transaction costs. The corresponding lower bounds for medium-variance stocks are 20.2% and 12.2%, respectively. In fact, the lower bounds presented in Table 3 are higher than the values obtained under either simplified tax trading or restricted tax trading in the absence of offset rules. This indicates that the ability to trade optimally at each date in the short-term region adds considerably to the value of the tax timing, even in the absence of the restarting option.

For those cases marked with an asterisk (*) in Table 3, the introduction of offset rules has no effect on the optimal realization policy or the value of tax timing. For those cases in which the optimal long-term realization boundary is close to $P/B = 1.0$ in the absence of offset rules, the lower bounds reported in Table 3 are fairly tight. For example, for medium-variance stocks with one-way transaction costs of $c = 0.5\%$, the optimal long-term realization boundary in the absence of offset rules is $P/B = 1.041$. For this case, the introduction of offset rules has almost no effect on the value of tax timing. For those cases in which the optimal long-term realization boundary is substantially higher than $P/B = 1.0$ in the absence of offset rules, the lower bounds reported in Table 3 may grossly underestimate the value of tax timing. For example, for high-variance stocks with zero transaction costs, the optimal policy is to realize all long-term capital gains each year in the absence of offset rules. While we do not expect this policy to survive in the presence of offset rules, we do expect that the optimal policy will involve the realization of some long-term capital gains late in the tax year when capital losses are small or nonexistent. To the extent that the restarting option has value to investors, the lower bounds reported in Table 3 will underestimate the total value of tax timing.

6. Concluding Comments

The results of this article suggest that the optimal trading of taxable securities in the presence of asymmetric capital gains taxes and transaction costs is more complicated than previously recognized. In the long-term region, investors optimally realize all long-term capital gains below some critical cutoff level. We derived an analytical expression for this optimal long-term cutoff level and explored its qualitative properties. We discovered that the optimal long-term real-

ization boundary is highly sensitive to the level of transaction costs, the volatility of the stock, the long-term and short-term tax rates, and the length of the short-term region. In the short-term region, the optimal trading policy involves the deferral of all capital gains and some smaller capital losses. Contrary to common belief, we showed that the deferral of short-term capital losses can be optimal even in the absence of transaction costs. The maximum short-term loss investors optimally defer is a function of the time remaining in the short-term region. For reasonable parameter values, we showed that it is not uncommon for investors to defer short-term capital losses as large as 10% in the absence of transaction costs.

Our results indicate that investors optimally realize nearly all long-term capital gains on medium- and high-variance stocks in the absence of transaction costs. However, we find that the benefits of realizing long-term capital gains to restart the short-term holding period are substantially reduced in the presence of relatively small transaction costs (see Tables 1 and 3). The benefits of restarting are further reduced in practice by the offset rules imposed by the actual tax code. Consequently, most investors are unlikely to find it optimal to realize long-term capital gains to restart the short-term holding period in practice. However, despite the relative insignificance of the restarting option in the presence of transaction costs and offset rules, the option to realize losses is largely unaffected by the introduction of offset rules and still retains considerable value in the presence of transaction costs. Our results indicate that for medium- and high-variance stocks, the ability to trade optimally in the short-term region adds an additional 12–30% to the stock's market value in the presence of 1% transaction costs (see Table 3).

The theoretical results discussed above are broadly consistent with the empirical results of Dammon, Dunn, and Spatt (1989). They find that while the option to realize losses is valuable to investors, even in the presence of transaction costs and offset rules, the option to realize long-term capital gains to restart the short-term holding period is generally of modest incremental value. What is puzzling about investors' observed trading behavior [see Poterba (1987) and Seyhun and Skinner (1994)] is the relative infrequency of capital loss realizations. This behavior may be the result of investors facing relatively high transaction costs, including the cost of monitoring portfolio positions, and/or the prevalence of portfolios composed primarily of assets with large embedded capital gains. Perhaps a general equilibrium model of investors' portfolio and consumption decisions in the presence of capital gains taxes and transaction costs can explain the observed trading behavior of investors. We leave this for future research.

Appendix: Derivation of the Valuation Equation for the Long-Term Region

To derive the valuation equation for the long-term region (i.e., for $b > N$), we begin by rearranging Equation (12) to give

$$v(x) - \pi_u v(ux) - \pi_d v(u^{-1}x) = [\pi_u u + \pi_d u^{-1}](1 - \tau_D)x, \quad (\text{A1})$$

where the dependence of v on b has been dropped for convenience. The general form of the solution to the above difference equation is

$$v(x) = A_1 x^m + A_2 x^n + A_3 x, \quad (\text{A2})$$

where $A_1 x^m$ and $A_2 x^n$ are solutions to the homogeneous equation

$$v(x) - \pi_u v(ux) - \pi_d v(u^{-1}x) = 0, \quad (\text{A3})$$

and $A_3 x$ is the particular solution to Equation (A1). Letting $v(x) = A_3 x$, Equation (A1) becomes

$$A_3 x[1 - \pi_u u - \pi_d u^{-1}] = [\pi_u u + \pi_d u^{-1}](1 - \tau_D)x. \quad (\text{A4})$$

Solving Equation (A4) for A_3 and using the definition of $\hat{\Pi}$ given by Equation (3) yields

$$A_3 = \hat{\Pi}(1 - \tau_D). \quad (\text{A5})$$

To find the two solutions to Equation (A3), let $v(x) = A_1 x^m = A_2 x^n$. Then Equation (A3) becomes

$$A_1 x^m[1 - \pi_u u^m - \pi_d u^{-m}] = A_2 x^n[1 - \pi_u u^n - \pi_d u^{-n}] = 0. \quad (\text{A6})$$

If $A_1 \neq 0$ and $A_2 \neq 0$, Equation (A6) requires

$$[1 - \pi_u u^m - \pi_d u^{-m}] = [1 - \pi_u u^n - \pi_d u^{-n}] = 0. \quad (\text{A7})$$

The solutions to Equation (A7) are

$$m = \{\ln[1 - (1 - 4\pi_u \pi_d)^{1/2}] - \ln[\pi_u] - \ln[2]\} / \ln[u] < 0 \quad (\text{A8})$$

and

$$n = \{\ln[1 + (1 - 4\pi_u \pi_d)^{1/2}] - \ln[\pi_u] - \ln[2]\} / \ln[u] > 1. \quad (\text{A9})$$

We next argue that because $n > 1$, A_2 must be equal to zero. The argument relies on the fact that there exists at most one optimal cutoff level in the long-term region, $0 < x_L < \infty$, below which all long-term capital gains and losses are realized and above which all long-term capital gains and losses are deferred. The proof that the optimal long-term realization boundary, x_L , is unique is provided in the lemma at the end of the appendix.

The function $v(x)$ is bounded above by the price (normalized by the basis), Πx , and below by 0. That is,

$$0 \leq v(x) \leq \Pi x. \quad (\text{A10})$$

Dividing Equation (A10) by $x > 0$ yields

$$0 \leq v(x)/x \leq \Pi, \quad (\text{A11})$$

where Π is a constant. According to Equation (A2) the value of $v(x)/x$ is

$$v(x)/x = A_1 x^{m-1} + A_2 x^{n-1} + A_3, \quad (\text{A12})$$

which is unbounded as $x \rightarrow \infty$ if $A_2 \neq 0$, since $n > 1$. Therefore, to satisfy Equation (A11), we require $A_2 = 0$. If $A_1 \neq 0$, then Equation (A12) is also unbounded as $x \rightarrow 0$, since $m < 0$. However, we do not require $A_1 = 0$, since the valuation equation $v(x)$ only applies for values of x in the deferral region (i.e., for values of x above the optimal long-term cutoff level, $x_L > 0$).²²

The constant A_1 is determined by imposing the following boundary condition on the solution:

$$v(x_L) = (1 - c)(1 - \tau_L)\Pi x_L + \tau_L. \quad (\text{A13})$$

Equation (A13) requires the value of deferral to equal the value of realizing at the optimal long-term realization boundary, x_L . Substituting $A_2 = 0$ into Equation (A12), Equations (A12) and (A13) together imply

$$A_1 [x_L]^m + A_3 x_L = (1 - c)(1 - \tau_L)\Pi x_L + \tau_L, \quad (\text{A14})$$

or, equivalently,

$$A_1 = [(1 - c)(1 - \tau_L)\Pi - (1 - \tau_D)\hat{\Pi}](x_L)^{1-m} + \tau_L(x_L)^{-m}. \quad (\text{A15})$$

²² The deferral region must extend above x_L rather than below it. If it were optimal to defer for all $x < x_L$, then to keep $v(x)/x$ from becoming unbounded as $x \rightarrow 0$ we would require $A_1 = 0$ and $v(x)$ would become

$$v(x) = A_2 x^n + A_3 x,$$

for all $x < x_L$. However, if it is optimal to defer for all $x < x_L$, then we also require

$$v(x) > (1 - c)(1 - \tau_L)\Pi x + \tau_L,$$

for all $x < x_L$. Letting $x = 0$, it is easy to see that the above two conditions are contradictory. Thus, the deferral region cannot extend below x_L , but must extend above it.

Therefore, the valuation equation is given by

$$v(x) = \{[(1-c)(1-\tau_L)\Pi - (1-\tau_D)\hat{\Pi}]x_L + \tau_L\}(x/x_L)^m + (1-\tau_D)\hat{\Pi}x, \quad (\text{A16})$$

for all $b > N$ and $x > x_L$. This completes the derivation.

Lemma. *The solution to Equations (12) and (13) has at most one optimal cutoff level, $0 < x_L < \infty$.*

Proof. The proof strategy is similar to that found in Williams (1985, p. 306). Assume there are two (or more) optimal long-term cutoff levels satisfying Equations (12) and (13). Let x_1 and x_2 , $0 < x_1 < x_2 < \infty$, denote two of these optimal long-term cutoff levels such that for all $x_1 < x < x_2$ it is optimal to defer the associated long-term capital gain or loss.²³ For all values of x in the deferral region, $x_1 < x < x_2$, the general solution for $v(x)$ is given by Equation (A2), where the coefficients m and n are given by Equations (A8) and (A9), respectively. At the optimal cutoff levels, x_1 and x_2 , we require the following two conditions to hold:

$$v(x_i) = (1-c)(1-\tau_L)\Pi x_i + \tau_L, \quad i = 1, 2 \quad (\text{A17})$$

and

$$v'(x_i) = (1-c)(1-\tau_L)\Pi, \quad i = 1, 2. \quad (\text{A18})$$

Equation (A17) is identical to Equation (13), and Equation (A18) is the familiar “tight-fit” condition that requires the marginal value of realizing at the optimal cutoff points, x_1 and x_2 , to be equal to the marginal after-tax, after-transaction costs proceeds from the sale. For all values of x in the deferral region, $x_1 < x < x_2$, the valuation equation must satisfy

$$v(x) > (1-c)(1-\tau_L)\Pi x + \tau_L. \quad (\text{A19})$$

Equations (A17), (A18), and (A19) imply that $v''(x_i) > 0$ for $i = 1, 2$. Consequently, there must also exist two inflection points, x_3 and x_4 , satisfying $v''(x_i) = 0$ for $i = 3, 4$, where $x_1 < x_3 < x_4 < x_2$. Differentiating Equation (A2) twice and multiplying through by x^2 , evaluated at x_3 and x_4 , yields

$$0 = m(m-1)A_1x_i^m + n(n-1)A_2x_i^n, \quad i = 3, 4. \quad (\text{A20})$$

²³ It is not possible for the deferral region to extend to all $x < x_1$ and $x > x_2$, with optimal realization for all $x_1 \leq x \leq x_2$. The proof is similar to that provided in the previous footnote.

Equation (A20) requires A_1 and A_2 to be of opposite sign since $m < 0$ and $n > 1$. Now subtract Equation (A20) evaluated at x_4 from Equation (A20) evaluated at x_3 :

$$0 = m(m-1)A_1(x_3^m - x_4^m) + n(n-1)A_2(x_3^n - x_4^n) \quad (\text{A21})$$

Since $x_3^m - x_4^m > 0$ and $x_3^n - x_4^n < 0$, Equation (A21) requires A_1 and A_2 to be of the *same* sign. Consequently, Equations (12) and (13) cannot hold at both x_1 and x_2 . This completes the proof. ■

References

- Bossaerts, P., and R. Dammon, 1994, "Tax-Induced Intertemporal Restrictions on Security Returns," *Journal of Finance*, 49, 1347-1371.
- Constantinides, G., 1983, "Capital Market Equilibrium with Personal Tax," *Econometrica*, 51, 611-636.
- Constantinides, G., 1984, "Optimal Stock Trading With Personal Taxes: Implications for Prices and the Abnormal January Returns," *Journal of Financial Economics*, 13, 65-89.
- Dammon, R., K. Dunn, and C. Spatt, 1989, "A Reexamination of the Value of Tax Options," *Review of Financial Studies*, 2, 341-372.
- Dammon, R., and C. Spatt, 1995, "Do Heterogeneous Basis Values Affect Equilibrium Asset Prices?," unpublished manuscript, Carnegie Mellon University.
- Dunn, K., and C. Spatt, 1986, "The Effect of Refinancing Costs and Market Imperfections on the Optimal Call Policy and the Pricing of Debt Contracts," unpublished manuscript, Carnegie Mellon University.
- Poterba, J., 1987, "How Burdensome are Capital Gains Taxes? Evidence From the United States," *Journal of Public Economics*, 33, 157-172.
- Schultz, P., 1988, *Tax Loss Selling in Common Stocks*, unpublished Ph.D. dissertation, University of Chicago.
- Seyhun, H. N., and D. Skinner, 1994, "How do Taxes Affect Investors' Stock Market Realizations? Evidence From Tax-Return Panel Data," *Journal of Business*, 67, 231-262.
- Williams, J., 1985, "Trading and Valuing Depreciable Assets," *Journal of Financial Economics*, 14, 283-308.