Should Investors Defer Long-Term Gains in Taxable Stock Portfolios?

Jeff Whitworth[1]

Abstract

Investors with taxable portfolios sometimes delay the sale of appreciated stock to defer capital gains taxes. While this strategy does help to reduce taxes, it can cause the portfolio to become more concentrated over time, leading to higher overall volatility and lower long-term returns. This paper evaluates the tradeoff between tax efficiency and diversification via Monte Carlo simulation and finds that diversification is far more important for the investor's terminal wealth, especially over longer time horizons. Under a reasonable set of assumptions, investors are better off rebalancing almost completely each year, even though it requires selling some recent winners and paying capital gains taxes. While tax efficiency does become somewhat more important for individuals who are taxed more heavily on gains or who expect an eventual step-up in basis, even these investors should tolerate only a modest increase in portfolio concentration.

Recommended Citation

Whitworth, J. (2024). Should investors defer long-term gains in taxable stock portfolios? *Financial Services Review, 32*(4), 13-26.

Introduction

How should taxable investors deal with appreciated stock in their portfolios? While owning stock that has increased in value is obviously a good "problem" to have, it does present investors with a dilemma. Selling would trigger a capital gains tax liability, so it may be advantageous to defer the realization of gains as long as possible. However, a policy of never selling appreciated shares inevitably leads to a loss of diversification over time, as the portfolio becomes increasingly dominated by a few big winners. The resulting increase in risk – while undesirable in itself – also reduces the portfolio's long-term expected growth rate through "volatility drag." Therefore, it is not immediately clear whether an investor should sell appreciated shares or continue to hold them.

In an early study on tax-efficient investing, Constantinides (1983) argues that it is best to realize losses as soon as they occur and to defer gains indefinitely (or until exogenous factors force the investor to liquidate shares). However, in a follow-up study, Constantinides (1984) demonstrates that when short-term gains are taxed as ordinary income but long-term gains are taxed at a lower rate, it often makes sense to realize some long-term gains so that future losses can be realized short-term and deducted against ordinary income. Smith and Smith (2008) propose a strategy of realizing all losses but also realizing enough gains to offset any losses in

[1] Corresponding author (WhitworthJ@uhcl.edu). University of Houston-Clear Lake, Houston, TX, USA

excess of $3,000 (the maximum that can be deducted against ordinary income in any year). They show via Monte Carlo simulation that this method generates higher terminal wealth than any of the strategies previously proposed by Constantinides.

In a more recent simulation study of tax-efficient investing, Whitworth (2018) considers the effects of loss recognition, Smith and Smith's (2008) limited gain recognition, and the use of appreciated stock for charitable donations. While all of these strategies clearly increase the investor's wealth over time, a more detailed breakdown of his results shows that most of the increase is attributable not to capital gains tax savings, but to improved portfolio diversification. This raises the question: Could diversification be so important that investors should actually sacrifice tax efficiency for a more balanced portfolio? If so, might investors be better off *completely* rebalancing their portfolios each year, even though it would require selling some appreciated shares and incurring capital gains taxes? Or could there be an optimal balance between tax efficiency and diversification that calls for less-than-complete annual rebalancing? These questions – which are not addressed in Whitworth (2018) – are the central focus of this paper.

It is well known that diversification reduces portfolio risk without reducing expected return – a fact sometimes referred to as the "only free lunch in investing." Since the long-term geometric growth rate on an asset is approximated by $g = \mu - 0.5\sigma^2$, where μ is the mean annual return and σ is the annualized return volatility, the reduction in portfolio risk that accompanies diversification should result in improved long-term growth. Indeed, several studies have noted the existence of a "diversification return" (Cuthbertson et al., 2016; Booth and Fama, 1992; Willenbrock, 2011; Erb and Harvey, 2006; Bouchey et al., 2012), and because of this, Feld (1999) suggests that many investors might be better off selling appreciated shares and taking the capital gains tax hit rather than dealing with the higher volatility inherent in an overly concentrated portfolio. In a simplified analysis of an investor who is fully concentrated in a single appreciated stock, Stein et al. (2000) find that "near-complete diversification" into a lower-volatility portfolio is preferable, "despite a high initial tax cost."

This paper considers the case of a taxable investor managing a multi-stock portfolio and simulates the likely long-term outcomes from different strategies. We assume that all losses are realized each year but also test several criteria for deciding when to sell winning stocks. Consistent with the studies cited above, our results indicate that the investor's after-tax terminal wealth is maximized by keeping the portfolio almost fully diversified via annual rebalancing, even though this method results in more capital gains taxes being paid. Even an investor with a high capital gains tax rate or one who expects a step-up in basis (which would increase the incentive to defer gains) should be willing to tolerate only a modest increase in portfolio concentration.

The next section describes the Monte Carlo simulation methodology used in this paper. The following section presents and analyzes the results of the study, and the final section concludes.

Simulation Methodology

This study builds on the simulation methodology of Whitworth (2018), which itself is based on Smith and Smith (2008). We assume that an investor has an initial portfolio of $250,000, consisting of $10,000 invested in each of 25 non-dividend-paying stocks. For each year t, a simulated market return R_{mt} is generated from a normal distribution with a mean of 8% and a standard deviation of 20%.[2] A set of idiosyncratic disturbances ε_{it} (i = 1, 2, … 25) is also generated via independent draws from a normal distribution with mean zero and (for reasons that will become clear shortly) a standard deviation of 34.641%. The set of realized stock returns R_{it} (i = 1, 2, …

[2] These market return assumptions are the same as those used in Whitworth (2018) and are generally consistent with the 2024 Long-Term Capital Market Assumptions of J.P. Morgan, which are available at https://am.jpmorgan.com/us/en/asset-management/institutional/insights/portfolio-insights/ltcma/.

25) in year t is then determined by adding the market return R_{mt} to each of the 25 disturbances ε_{it}. Given the parameters above, the simulated annual stock returns R_{it} will be distributed normally with a mean of 8% and a standard deviation of

$$\sigma_i = \sqrt{\sigma_m^2 + \sigma_\varepsilon^2} = \sqrt{0.2^2 + 0.34641^2} = 0.4 = 40\%.$$ [3]

The contemporaneous correlation between the returns of any two different stocks i and j in the same year will be

$$\rho_{ij} = \frac{\sigma_{ij}}{\sigma_i \sigma_j} = \frac{\sigma_m^2}{\sigma_i \sigma_j} = \frac{0.2^2}{0.4 \cdot 0.4} = 0.25.$$

These values for σ_i and ρ_{ij} are the same as those used by Smith and Smith (2008) and Whitworth (2018). Returns across *different* years are generated independently and thus are uncorrelated.[4] Ordinary income is taxed at 24%, while long-term capital gains are taxed at 15%[5].

Given these assumptions, the investor does not try to time the market. Of course, one could adopt a buy-and-hold approach where no stocks are sold until the end of the investment horizon, but we consider several alternative strategies (described below) where trading decisions are made based on taxes and portfolio diversification. At the end of each year, any lot of stock whose value has declined below its original basis is sold to realize capital losses. Any capital gains that have accrued are then handled according to one of three policies. The first two policies are considered in Smith and Smith (2008) and in Whitworth (2018); however, the third policy has not been considered in previous simulation studies. The three strategies for dealing with accrued gains that we examine are as follows:

1) *Defer All Gains*: Appreciated shares are held as long as possible (i.e., until the end of the investment horizon) so long as their value never dips below the original basis.
2) *Realize Gains to Offset Excess Losses*: Under this policy, the investor follows Smith and Smith's (2008) strategy of realizing only enough gains to offset any losses exceeding $3,000 (the maximum which may be deducted against ordinary income in any year). Because an important objective of this strategy is to facilitate rebalancing, shares are sold first from the portfolio's most heavily-weighted stock until its remaining value equals that of the second heaviest-weighted. Then shares are sold from the first two stocks until each of their values equal that of the third heaviest-weighted. This continues until the desired amount of capital gains have been realized, or until there are no more unrealized gains left in the portfolio.
3) *Rebalance the Entire Portfolio*: Under this policy, if any stock comprises more than 4% (i.e., 1/25) of the portfolio, enough shares are sold to bring that stock's weight down to 4%.

When selling shares under policy (2) or (3) to reduce a stock's weight in the portfolio, it is possible that the investor might own multiple lots of that company's stock. If so, shares are sold first from the lot with the highest basis-to-value ratio.

Under policies (1) and (2) – and possibly under policy (3) – the investor will realize a net capital

[3] As noted in Whitworth (2018), this value for σ_i is reasonable given the individual stock volatilities that have been observed historically [e.g., see Statman (1987) and Campbell et al. (2001)]. To the extent that investors choose lower-volatility stocks, less rebalancing would be needed to maintain diversification, and tax efficiency would likely become more of a consideration. (Of course, the reverse is true if investors choose higher-volatility stocks.)

[4] The assumption of serially independent returns is largely consistent with the market efficiency literature summarized by Fama (1970, 1991) and Malkiel (2003). To the extent that stock returns exhibit momentum, Constantinides' (1983) strategy of recognizing losses and continuing to hold winners would increase in effectiveness. To the extent that returns are mean-reverting, regular portfolio rebalancing (which requires selling recent winners and buying additional shares of recent losers) would become more desirable.

[5] According to the most recently available IRS Statistics of Income (for Tax Year 2021), 67% of returns reporting long-term gains were subject to the 15% tax rate. As of 2024, this rate applies to single filers with taxable income between $47,026 and $518,900 and to married couples filing jointly with income between $94,051 and $583,750.

loss[6], up to $3,000 of which will be deducted against ordinary income (resulting in a tax savings of 24% of the deduction). Any unused loss deduction is carried forward to the next year. However, policy (3) will often realize more capital gains than losses, in which case the net gain is taxed at the long-term rate of 15%. (This assumes that the appreciated shares were held for at least a year plus one day to qualify for the long-term capital gains rate.)

The proceeds from end-of-year stock sales (plus the tax savings on net losses, or minus taxes paid on net gains) are then reallocated to the least-weighted stocks in the portfolio. Funds are invested first in the least-weighted stock until its weight equals the second least-weighted; then funds are invested in the two least-weighted stocks until their respective weights match the third; and so on, until all cash has been reinvested. This rebalancing procedure ensures that the portfolio is as diversified as possible.[7]

This process continues each year until the end of the investment horizon (assumed to be 10, 20, 30, 40, or 50 years), at which point the entire portfolio is liquidated and taxes are paid on the net long-term gain.

For each iteration of the simulation, we concurrently track the value of a buy-and-hold portfolio which begins with $250,000 invested equally across the same 25 stocks, but in which no stocks are ever sold until the final liquidation date. From this, we calculate the percentage by which the investor's after-tax terminal wealth exceeds that of the buy-and-hold portfolio. Some of this wealth increase may be due to improved tax efficiency, while some is attributable to the "diversification return" previously noted. To determine how much is due to better diversification, we also track the value of a portfolio which – like the buy-and-hold portfolio above – starts with the same $250,000 and never realizes any gains or losses until the terminal date. In this portfolio, however, funds are costlessly redistributed each year (without any immediate tax consequence) so that the individual stock weights match those in the investor's portfolio. Therefore, this portfolio receives none of the tax benefits that the investor does via tax-efficient trading, but it maintains the same level of diversification. By also comparing its performance versus buy-and-hold, we can ascertain how much of the investor's wealth gain is due to diversification and how much is from tax savings.

For each of the three gain realization policies above and for each possible time horizon (ranging from 10 to 50 years), one million iterations of the simulation are run to obtain parameter estimates.

Results

Effect of the three gain realization strategies on terminal wealth and risk

Table 1 shows the mean and median increases in after-tax terminal wealth (relative to a buy-and-hold approach) from implementing the three investment strategies described in the previous section. In addition, the table shows how much of the increase is attributable to tax savings and how much is due to improved diversification. It is immediately clear that all three strategies outperform buy-and-hold and that the difference is greater over longer time horizons. A comparison of Panels A and B shows that Smith and Smith's (2008) strategy of realizing enough gains to offset excess losses does much better than a policy of simply deferring all gains. The middle column of the table shows that very little of the wealth increase is due to additional tax savings. This is not surprising since recognizing gains is generally not advantageous from a tax standpoint, except that effectively resetting the cost basis on part of the portfolio may create a few more opportunities to harvest future losses. However, as seen in the rightmost column of the table, there is a large diversification return. This

[6] In the unlikely event that none of the 25 stocks in the portfolio experienced a loss, no gains or losses would be realized that year under policies (1) and (2).

[7] Implementing this strategy usually will require repurchasing shares that have just been sold to recognize capital losses. In practice, this would run afoul of wash sale rules. However, an investor can circumvent these restrictions by waiting 31 days to repurchase the shares, or by immediately purchasing shares of a different stock with similar risk and return characteristics.

is also as expected since the main benefit of Smith and Smith's approach is that it routinely rebalances the portfolio (albeit to a limited extent) by harvesting gains from the most heavily-weighted stocks and reinvesting the funds into the least-weighted stocks. This in turn reduces the volatility of the portfolio and the consequent "volatility drag" on returns.

Table 1. Mean (Median) Increase in After-Tax Terminal Wealth from Alternative Investment Strategies vs. Buy-and-Hold

Years Invested	Total Increase	From Tax Savings	From Diversification
	Panel A: Realize Losses and Defer All Gains		
10	1% (1%)	1% (1%)	0% (0%)
20	4% (2%)	2% (2%)	0% (0%)
30	8% (3%)	4% (3%)	4% (0%)
40	13% (4%)	4% (3%)	9% (1%)
50	22% (6%)	7% (4%)	15% (2%)
	Panel B: Realize Gains to Offset Excess Losses		
10	3% (2%)	1% (1%)	2% (1%)
20	11% (8%)	2% (2%)	9% (6%)
30	26% (17%)	4% (3%)	22% (14%)
40	48% (27%)	6% (4%)	42% (23%)
50	78% (39%)	9% (5%)	69% (34%)
	Panel C: Rebalance Entire Portfolio		
10	3% (3%)	1% (1%)	2% (2%)
20	12% (10%)	0% (0%)	12% (10%)
30	32% (24%)	0% (-1%)	32% (25%)
40	64% (44%)	-3% (-4%)	67% (48%)
50	116% (75%)	-9% (-9%)	125% (84%)

Each investment strategy is simulated 1,000,000 times over 10, 20, 30, 40, or 50 years. The simulated portfolio begins with $250,000 invested equally across 25 non-dividend-paying stocks. Each year, we generate a set of 25 stock returns with mean 8%, standard deviation 40%, same-year inter-stock correlation 0.25, and no intertemporal correlation. Capital losses are always realized annually. The three respective panels report results assuming that (a) no capital gains are ever realized until the end of the timeline, (b) enough gains are realized each year to offset any losses in excess of $3,000, or (c) the portfolio is completely rebalanced each year. Net realized losses are deducted (up to a maximum of $3,000 per year) against ordinary income, which is taxed at 24%. Net realized gains are taxed at 15%. After-tax proceeds from stock sales are reinvested in the least-weighted stocks. In each iteration the final after-tax wealth is compared to that of (1) a portfolio using a pure buy-and-hold approach, and (2) a control portfolio that remains untaxed until the end of the horizon but is costlessly rebalanced every year so that its individual stock weights match those in the simulated portfolio. The total increase column reports the mean (and median) percentage by which the simulated portfolio's terminal value exceeds that of the buy-and-hold portfolio. The tax savings component is computed by comparing the simulated portfolio to the control portfolio that is costlessly rebalanced each year. The diversification component is the total increase minus the tax savings component.

Since the added diversification of Smith and Smith's strategy provides such large benefits, a natural question is whether an investor might want to diversify even further – perhaps even rebalancing the whole portfolio every year – even though doing so will require realizing and paying taxes on some gains. Looking at Panel C, the answer is a clear "yes." Even though (as expected) the tax impact of this strategy is worse than the others and actually affects wealth negatively at longer horizons, the additional diversification gains more than make up for this. The net result is that fully rebalancing each year outperforms Smith and Smith's more limited gain recognition strategy.

Table 2 shows the effect of the three strategies on risk, as measured by the standard deviation of terminal wealth in our simulation. As seen in the first column of the table, realizing only losses has almost no effect on portfolio risk, as the standard deviation of terminal wealth is very near what it would be under a buy-and-hold strategy. Smith and Smith's (2008) strategy does reduce risk modestly. However, a policy of full portfolio rebalancing does much better, especially over longer horizons. Over a 50-year time frame, the standard deviation of terminal wealth for someone using the more limited gain harvesting approach would be about 80% of a buy-and-hold investor's risk; however, for an investor who fully rebalances each year, it would be only half of that (i.e., 40% of the buy-and-hold risk). The difference between full rebalancing and Smith and Smith's more limited gain recognition is relatively modest over 10- or 20-year horizons, but it becomes much larger over 40-50 years.

Table 2. Ratio of Standard Deviation of Terminal Wealth from Alternative Investment Strategies vs. Buy-and-Hold

Years Invested	Realize Losses, Defer All Gains	Realize Offsetting Gains	Rebalance Entire Portfolio
10	1.01	0.98	0.94
20	1.01	0.95	0.83
30	1.01	0.91	0.68
40	1.01	0.87	0.52
50	0.99	0.80	0.40

The simulation procedure is described in Table 1. For each investment strategy and time horizon, this table reports the standard deviation of the 1,000,000 simulated portfolio terminal values divided by the standard deviation of the 1,000,000 corresponding buy-and-hold portfolio values.

When comparing Tables 1 and 2, it is clear that gain harvesting is most effective at the longest horizons, both at reducing risk and at improving mean and median terminal wealth. As Table 1 shows, the diversification return (which is a direct consequence of the lower risk) becomes significantly greater over time. This is not surprising because these longer horizons are precisely where portfolios left to themselves could otherwise become extremely unbalanced. To see why these gain harvesting strategies are especially effective in the long run, it is informative to look directly at how they impact the portfolio's level of diversification.

Table 3 shows how concentrated a portfolio may become in any one stock over different time frames. Concentration is measured here by the percentage of the portfolio's value that is invested in its heaviest-weighted stock. (This is a reasonable proxy for more sophisticated diversification measures – for example, a Herfindahl-style measure equal to the sum of the squares of the individual stock weights.) Bessembinder (2018) notes that "the compounding of random returns over multiple periods will typically impart positive skewness to longer horizon returns, even if the distribution of single-period returns is symmetric." As a consequence of this return skewness, we would

expect a portfolio left to itself to become increasingly unbalanced over time, as the results in the first column of Table 3 confirm. Under a buy and hold strategy, just one of the 25 stocks will on average make up 20% of the portfolio (which is not insignificant) after only 10 years. After 50 years, a single stock is likely to become more than half (53%) of the portfolio. Merely realizing losses helps a little bit, as the second column of the table shows slightly reduced maximum weights versus the first column. Smith and Smith's (2008) strategy (corresponding to the third column) does even better, but it still allows for one stock to become about one-third of the portfolio over longer horizons. By design, a strategy of always rebalancing the whole portfolio keeps each of the 25 stock weights at exactly 4%, which is optimal from a risk perspective. As a comparison of the third and fourth columns shows, this is a little better diversified after 10 years than a portfolio managed according to Smith and Smith's strategy, but after 40 or 50 years, it is drastically better. This explains why (as previously seen in Tables 1-2), the difference between the two strategies with respect to terminal wealth and risk is relatively small for shorter horizons but much greater over longer horizons.

Table 3. Mean Portfolio Weight of Largest Stock in Final Year

Years Invested	Buy and Hold	Realize Losses, Defer All Gains	Realize Offsetting Gains	Rebalance Entire Portfolio
10	20%	19%	8%	4%
20	32%	30%	14%	4%
30	41%	38%	22%	4%
40	48%	43%	29%	4%
50	53%	48%	35%	4%

The simulation procedure is described in Table 1. Results in the four respective columns are reported assuming that (a) no stocks are sold until the end of the horizon; (b) losses are realized annually but no gains are realized until the end of the horizon; (c) losses are realized annually, and enough gains are realized to offset any losses in excess of $3,000; or (d) losses are realized annually, and enough gains are realized to fully rebalance the portfolio annually. For each strategy and time horizon considered, this table reports the mean percentage of the portfolio that is invested in its most heavily-weighted stock at the end of the horizon.

Effects of Less than Complete Diversification

In the previous subsection, we saw that the best results (in terms of terminal wealth, risk, and diversification) were obtained by disregarding some of the tax-efficient strategies previously proposed in the literature and simply diversifying the portfolio fully every year instead. However, the wealth-maximizing policy[8] still might lie somewhere between full diversification and the Smith and Smith (2008) strategy (which still allows for considerable portfolio concentration, as seen in Table 3).

We now consider the effects of modifying policy (3) described earlier by relaxing the full diversification criterion to varying extents. Under this modification, the investor sets a maximum portfolio weight which no individual stock may exceed. This "weight limit" (which remains the same for the entire investment period) may be set as low as 4% if full annual rebalancing is desired, or at a higher percentage if the investor wishes to defer more capital gains while still not allowing the portfolio to get too far out of balance. At the end of each year, if any stock exceeds the weight limit, enough shares are sold to bring its weight

[8] In the analysis that follows, we focus mostly on how different strategies affect median terminal wealth because medians are less influenced by outliers than means. Although the graphs are not shown here, we also find that mean terminal wealth is affected similarly.

in the portfolio down to the limit. All losses are still recognized, just as in the original strategy. We repeat the simulation for weight limits from 4% to 20% (in increments of 1%).

Figure 1. Median Increase in After-Tax Terminal Wealth vs. Buy-and-Hold for Alternative Portfolio Weight Limits

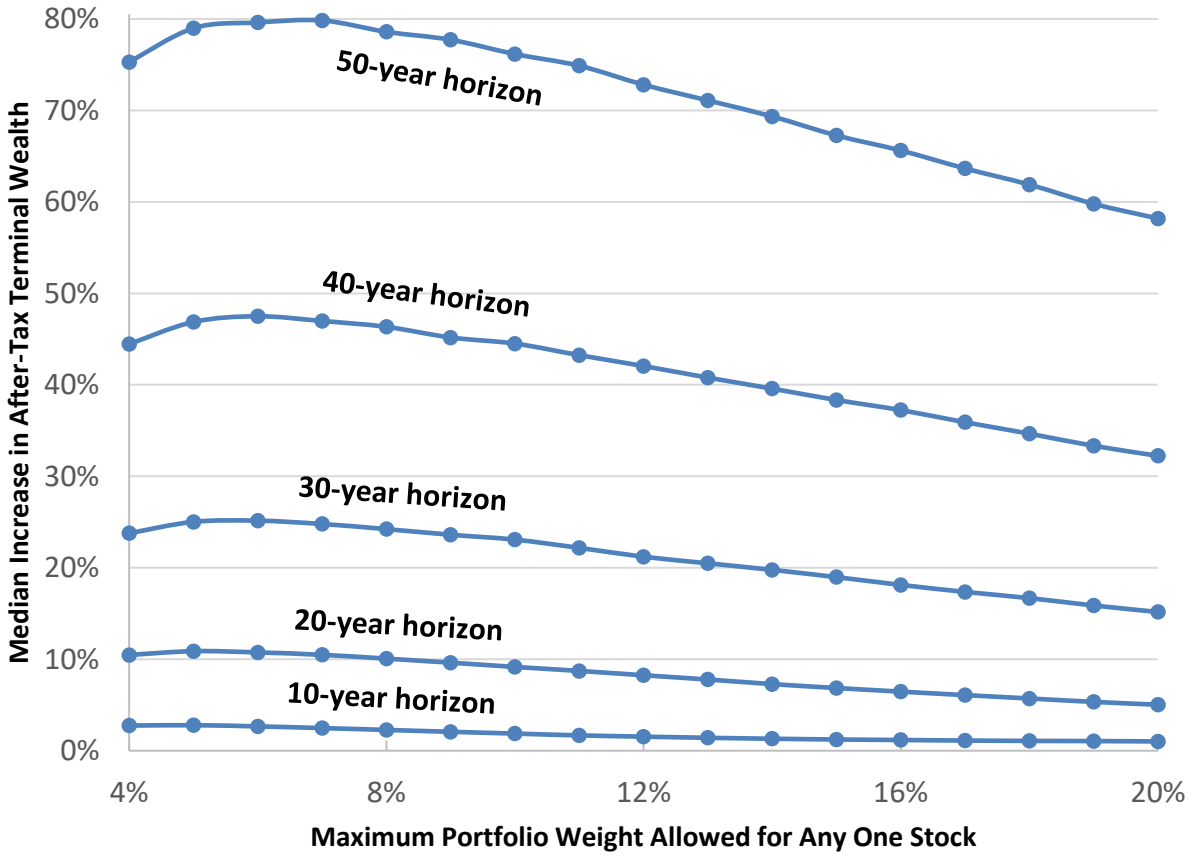

Each simulated portfolio begins with $250,000 invested equally across 25 non-dividend-paying stocks. Each year, we generate a set of 25 stock returns with mean 8%, standard deviation 40%, contemporaneous correlation 0.25, and no intertemporal correlation. Capital losses are always realized annually. A maximum weight is defined which no stock may exceed in the portfolio; if any stock's weight is over the limit at the end of a year, enough shares are sold to bring its weight down to the limit. Net capital losses (up to $3,000 per year) are deducted against ordinary income (which is taxed at 24%), and net capital gains are taxed at 15%. After-tax proceeds from stock sales are reinvested in the portfolio's least-weighted stocks. The final after-tax wealth at the end of the horizon is compared to a buy-and-hold portfolio over the same period. The simulation is repeated 1,000,000 times for investment horizons from 10 to 50 years and for maximum stock weights from 4% to 20%. For each horizon and maximum weight, this graph shows the median percentage by which the simulated portfolio's terminal value exceeds that of the buy-and-hold portfolio.

Figure 1 shows the median percentage increase in after-tax terminal wealth (again, relative to a buy-and-hold strategy) for different portfolio weight limits that an investor may set. For 10- and 20-year horizons, a 5% weight limit is optimal, although there is less variation between the outcomes of the alternative strategies over these shorter time frames. For 30- or 40-year horizons, a 6% weight limit does slightly better. Over 50 years, median terminal wealth is maximized by setting a 7% weight limit. Beyond that point, terminal wealth steadily declines as the portfolio

weight limit is raised, and the differences in wealth are especially notable over longer horizons. Given the assumptions of our simulation, near-complete diversification is best. It should be noted here that "best" refers only to the policy that maximizes median terminal wealth. In practice, investors may be willing to sacrifice some wealth in exchange for lower portfolio volatility. Although we do not attempt here to explicitly model the investor's utility as a function of risk and expected return, we acknowledge that in many cases the portfolio "weight limit" that maximizes utility may be slightly lower than the one that maximizes expected wealth. For example, in the preceding example, instead of setting the weight limit at 7%, an investor with a 50-year horizon would almost certainly prefer to set it at 6% or even 5% because doing so will reduce portfolio risk while sacrificing only a negligible amount of median terminal wealth. (It is less clear whether they would lower it further to 4% since that would reduce median wealth more discernibly.) Nevertheless, median terminal wealth is a key measure of the effectiveness of any investment strategy, so it is a reasonable starting point for the discussion.

Portfolio risk (as measured by the standard deviation of terminal wealth) is an increasing function of the weight limit since a higher weight limit leads to greater portfolio concentration and less diversification over time. However, it is not immediately clear whether the weight limit is always positively related to *shortfall* risk (i.e., the risk of earning less than a predetermined rate of return). Figure 2 shows the probability of ending up with a negative overall return for different strategies. Although the differences are relatively modest across alternative weight limits, shortfall risk is still a monotonically increasing function of portfolio concentration. For each horizon length, the risk is lowest with full diversification and steadily increases as higher portfolio weights are allowed.

Figure 2. Risk of Negative Overall Return for Alternative Portfolio Weight Limits

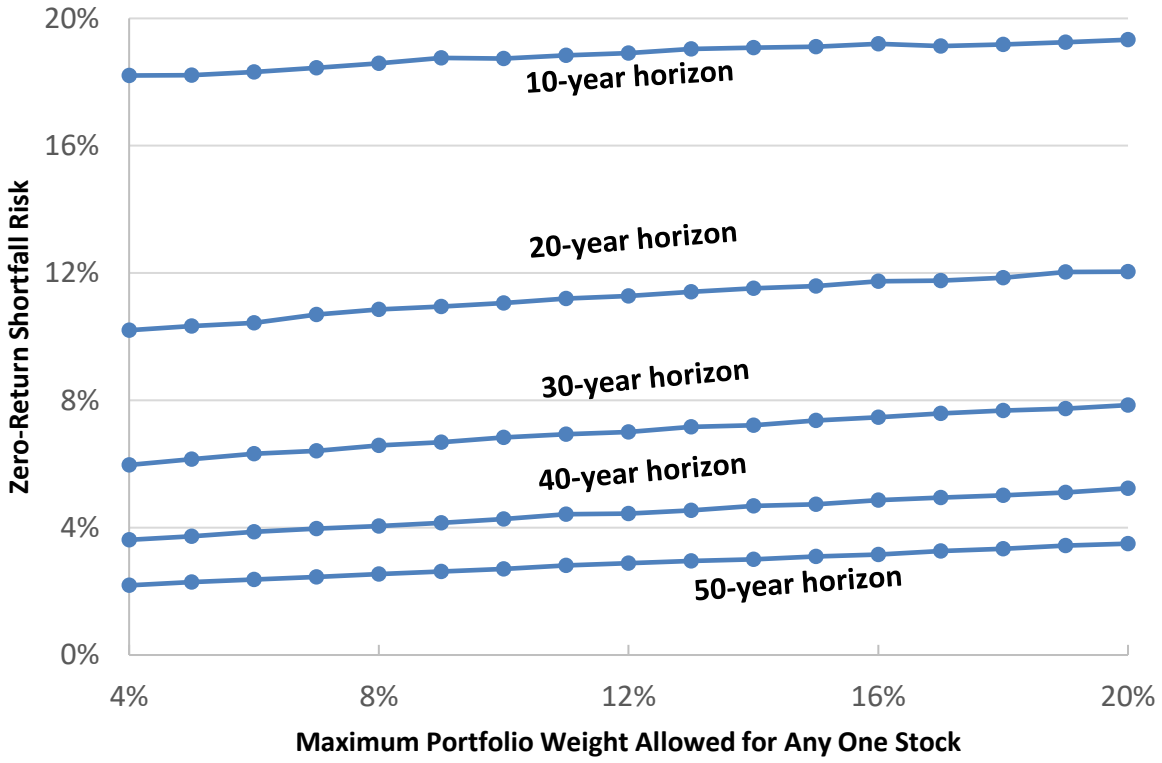

The simulation procedure is described in Figure 1. For each horizon and maximum stock weight, this graph shows the percentage of outcomes for which the simulated portfolio's terminal value is less than the initial $250,000 investment.

If Capital Gains are Untaxed on the Terminal Date

So far, we have assumed that all unrealized capital gains are finally taxed when the portfolio is liquidated at the end of the investment horizon. This makes sense if we interpret the terminal date as the point when the individual intends to consume the wealth they have accumulated. However, other interpretations are possible. For example, if the individual does not intend to consume the assets in the portfolio but instead plans to pass them on to heirs, the terminal date might be interpreted as the investor's death, at which point the portfolio's cost basis will be stepped up to current market value, erasing all accrued capital gains tax liability for the inheritors. If so, then it may make sense in some situations to continue holding appreciated stocks despite the increase in portfolio concentration. We would not expect the investor to completely ignore diversification concerns, of course, but the anticipation of a step-up in basis may increase the "weight limit" they would allow.

In Figure 3, we see that this is indeed true. When capital gains are untaxed on the terminal date, the wealth-maximizing weight limit shifts slightly to the right. For all horizons, median terminal wealth is maximized with a 7% or 8% weight limit (and in all cases, the wealth difference at 7% versus 8% is so slight as to be almost indistinguishable).

In reality, most investors probably do not have a single "terminal date" marking the end of the investment horizon but will likely consume some of their wealth over time (resulting in some taxable redemptions) before eventually passing the remainder on to heirs (who will enjoy a stepped-up basis on that portion of the portfolio). It is beyond the scope of this study to model all of these possible complexities, but it is reasonable to believe that in practice the optimal "weight limit" may lie somewhere between the values derived from Figures 1 and 3. It is also possible that one might use a dynamic policy that changes over time. For example, an individual who is nearing the end of life and expecting a stepped-up basis relatively soon might almost always defer gains, whereas someone with many years to go might realize more gains to keep the portfolio better diversified.

If Capital Gains are Taxed at a Higher Rate

It is important to note that the conclusions we have drawn so far depend on the particular parameters of this simulation. One important factor that may influence an investor's willingness to rebalance the portfolio is the rate at which any realized gains would be taxed. Under current law, long-term capital gains are generally taxed at substantially lower rates than ordinary income is. However, it is not guaranteed that these preferentially low rates will always exist in the future. From time to time, policymakers have proposed taxing long-term capital gains at the higher ordinary income tax rate. We expect that such a policy (or more generally, any increase in the capital gains tax rate) would increase the cost of portfolio rebalancing and therefore increase the degree of portfolio concentration that an investor would tolerate. This would be especially true if the investor is expecting much or all of the portfolio to be stepped up in basis at the end of the investment horizon.

Figure 3. Median Increase in Untaxed Terminal Wealth vs. Buy-and-Hold for Alternative Portfolio Weight Limits

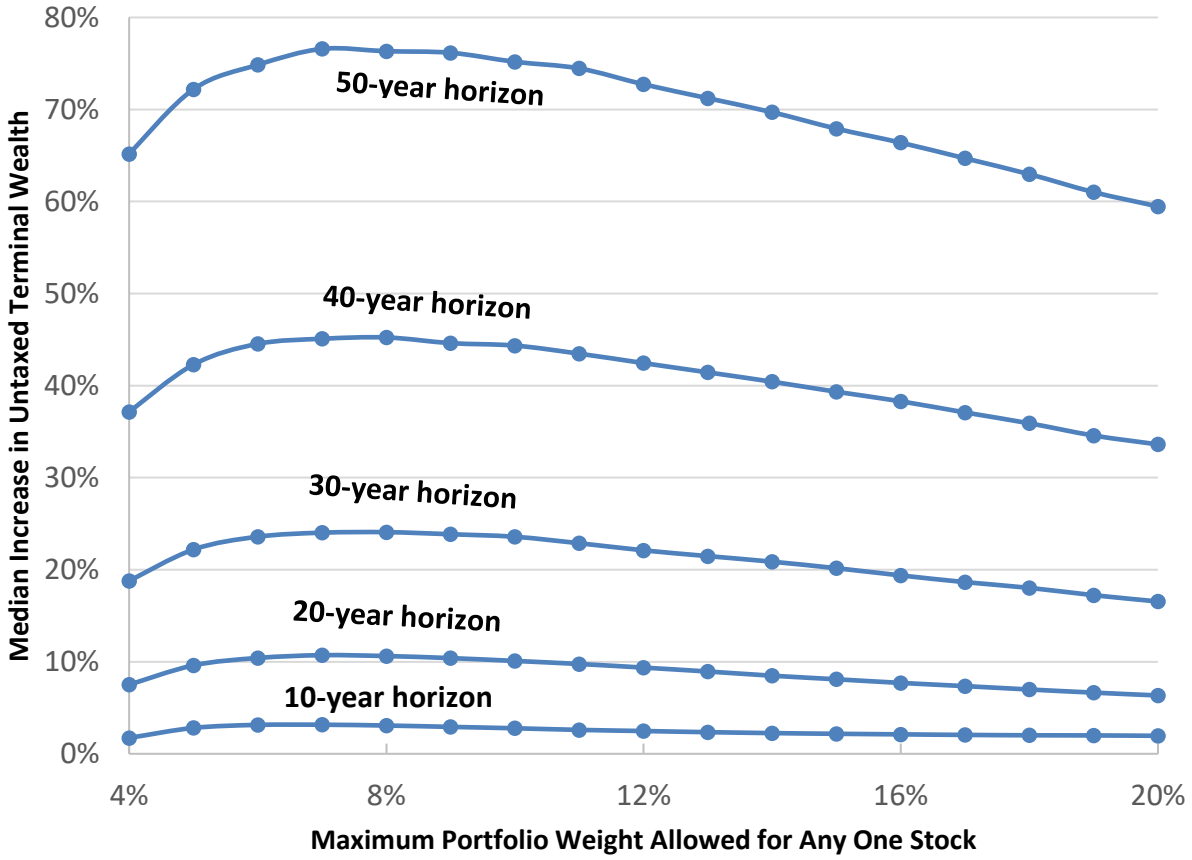

The simulation procedure is as described in Figure 1, except that any unrealized capital gains at the end of the investment horizon are not taxed. After each iteration, the final wealth is compared to a buy-and-hold portfolio which is also untaxed at the end of the horizon. For each horizon and maximum weight, this graph shows the median percentage by which the simulated portfolio's terminal value exceeds that of the buy-and-hold portfolio.

To see this, we repeat the previous simulation but assume that capital gains and ordinary income are both taxed at 40%, which is close to the current top individual income tax rate. The resulting median increases in terminal wealth relative to buy-and-hold are shown in Figure 4. As expected, the wealth-maximizing weight limit is greater when realized capital gains are taxed more heavily. For the 10- or 20-year horizons, median wealth is maximized when the investor sets a weight limit of 9% or 11%, respectively (although at these horizons, the wealth differences at 9% versus 11% are virtually indistinguishable). For horizons of 30 years or longer, the wealth-maximizing weight limit is 12%. From this it is clear that an increase in the capital gains tax rate increases the degree of portfolio concentration an investor should be willing to tolerate since there is a higher tax liability created by the transactions associated with rebalancing. In fact, it is interesting to note that under this scenario, full annual rebalancing can sometimes be slightly *worse* than buy-and-hold, as evidenced by the negative wealth increase when setting a 4% weight limit for some of the shorter time frames. This is consistent with the increased tax cost of rebalancing transactions and the large number of such transactions that would be required to keep the portfolio equally weighted over time. However, even with the

higher tax cost, the portfolio still should not be allowed to get too far out of balance if the objective is to maximize terminal wealth.

Transaction Costs

Unfortunately for investors, rebalancing one's portfolio regularly also means incurring transaction costs regularly. As capital gains taxes do, brokerage commissions and bid-ask spreads add to the cost of rebalancing and should serve to increase the portfolio concentration that an investor is willing to tolerate. But are transaction costs large enough to make a substantial difference in the wealth-maximizing strategy?

We investigate this question by repeating the previous simulation (where realized capital gains are taxed at 40%), except that every sale incurs a cost of 25 basis points (i.e., 0.25% of the value of the shares sold). Transaction costs obviously can be higher or lower depending on the characteristics of the stock and/or the trading environment, but this is a reasonable approximation of the effective transaction costs actually observed in U.S. stock markets (e.g., Hasbrouck, 2009; Chen and Velikov, 2023). In unreported results, we find that our simulated terminal wealth values are almost identical to those reported in Figure 4. In other words, transaction costs in the major U.S. equity markets are low enough that they do not substantially influence the investor's decision about whether to defer capital gains or recognize some gains to rebalance the portfolio.

Figure 4. Median Increase in Untaxed Terminal Wealth vs. Buy-and-Hold for Alternative Portfolio Weight Limits (Tax Rate = 40%)

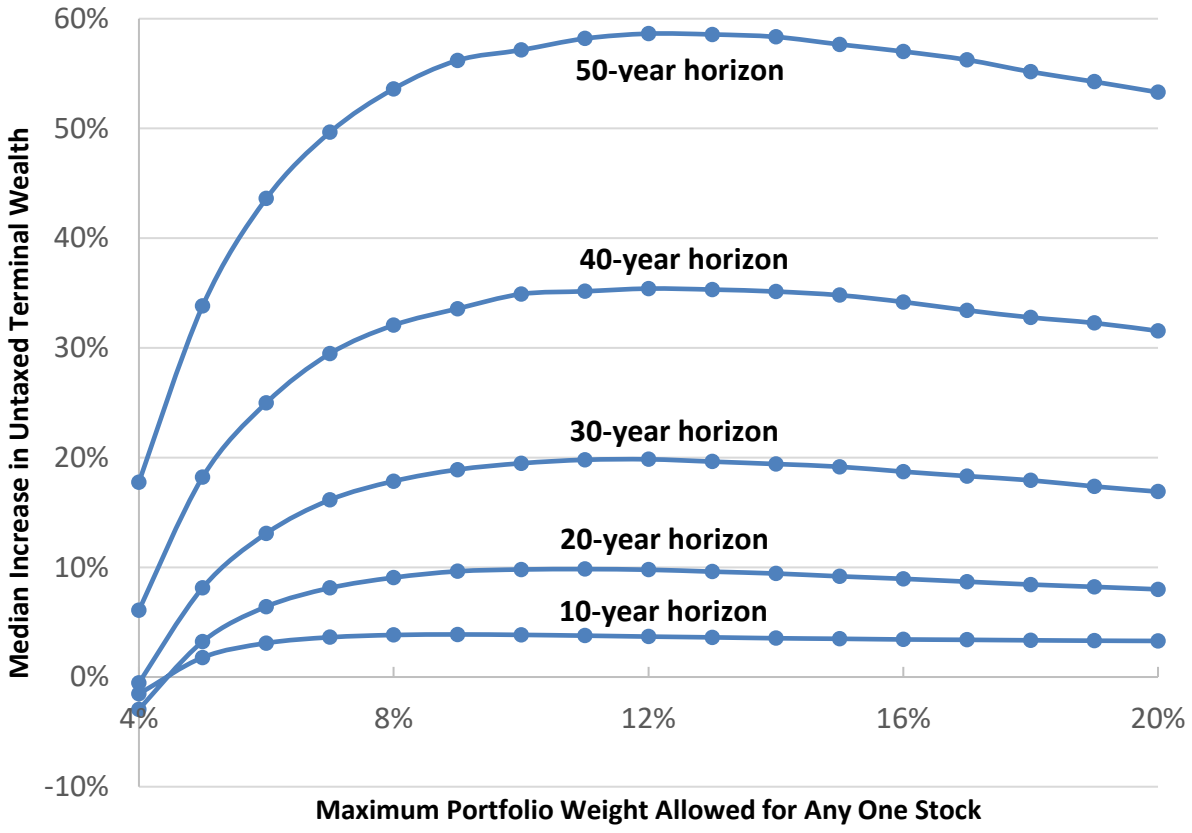

The simulation procedure is as described in Figure 3, except that ordinary income and any net realized capital gains are taxed at 40%. Unrealized gains at the end of the investment horizon remained untaxed. For each horizon and maximum weight, this graph shows the median percentage by which the simulated portfolio's terminal value exceeds that of the buy-and-hold portfolio.

Other Considerations

While it is beyond the scope of this study to examine every possible scenario, it is appropriate to acknowledge how the study's design and some of its specific assumptions may impact the results. First, our conclusions are based on a set of simulated stock returns generated using the methodology and parameters previously described. While this is reasonable and consistent with prior literature (e.g., Smith and Smith, 2008; Whitworth, 2018), other studies have employed a historical resampling approach when investigating similar questions, and it is possible that one might obtain different results by using a different methodology.

Regarding the specifics of our simulation, one simplifying assumption made is that the stocks do not pay dividends. All other things being equal, taxable individual investors generally should avoid dividend-paying stocks since dividends are tax-inefficient. Receiving dividends essentially converts capital gains which could be taxed much later (or not at all, if the investor expects a step-up in basis) into current income which is taxed immediately. In addition, the *entire* dividend distribution is taxed as income, unlike shares which are liquidated for rebalancing purposes (where only the value above the cost basis is taxed). Nevertheless, while some investors do own dividend-paying stocks, most dividends are relatively small and unlikely to have a significant impact on the investor's terminal wealth.

Another simplifying assumption of the simulation is that the investor makes no new contributions to the portfolio on an annual basis. In cases where one does make regular additional contributions (e.g., from wage income), the newly invested funds can be allocated to the least-weighted stocks, thereby improving portfolio diversification and reducing (but not eliminating) the need to harvest gains for the purpose of rebalancing. It may be of interest for future studies to simulate the outcomes of alternative strategies when investments are made regularly over time, rather than as a lump sum at the beginning.

Conclusion

This paper has considered alternative policies for when individuals with taxable portfolios should realize capital gains on their winning stocks. Constantinides (1983) advocates holding winners as long as possible, a strategy which is tax-efficient but can allow the portfolio to become very unbalanced over time. Smith and Smith (2008) propose a limited gain recognition strategy which does not hinder tax efficiency but does allow for greater portfolio diversification. We show that this leads to much higher terminal wealth, and that almost all of this is due to the "diversification return." We then show that the investor can do better still by rebalancing the portfolio almost completely each year. Even though this strategy requires paying some capital gains taxes earlier than if all gains were deferred, the increased diversification return overwhelms the relatively small reduction in tax efficiency. This, combined with the fact that lower risk is desirable in itself, suggests that taxable stock investors might be better off foregoing the usual advice to delay the sale of appreciated stocks as long as possible. Individuals who are taxed more heavily on capital gains or who expect an eventual step-up in basis may wish to defer more gains than a typical investor would. However, even in these cases, only a modest increase in portfolio concentration should be considered since diversification has such a strong impact on long-term returns.

References

Bessembinder, H. (2018). Do stocks outperform Treasury bills? *Journal of Financial Economics*, 129(3), 440-457. https://doi.org/10.1016/j.jfineco.2018.06.004

Booth, D. G. and Fama, E. F. (1992). Diversification returns and asset contributions. *Financial Analysts Journal*, 48(3), 26-32. https://doi.org/10.2469/faj.v48.n3.26

Bouchey, P., Nemtchinov, V., Paulsen, A., and Stein, D.M. (2012). Volatility harvesting: Why does diversifying and rebalancing create portfolio growth? *Journal of Wealth Management*, 15(2), 26-35.

Campbell, J. Y., Lettau, M., Malkiel, B. G., and Xu, Y. (2001). Have individual stocks become more volatile? An empirical exploration of idiosyncratic risk. *Journal of Finance*, 56(1), 1-43. https://doi.org/10.1111/0022-1082.00318

Chen, A. Y. and Velikov, M. (2023). Zeroing in on the expected returns of anomalies. *Journal of Financial and Quantitative Analysis*, 58(3), 968-1004. https://doi.org/10.1017/S0022109022000874

Constantinides, G. M. (1983). Capital market equilibrium with personal tax. *Econometrica*, 51(3), 611-636. https://doi.org/10.2307/1912150

Constantinides, G. M. (1984). Optimal stock trading with personal taxes: Implication for prices and the abnormal January return. *Journal of Financial Economics*, 13(1), 65-89. https://doi.org/10.1016/0304-405X(84)90032-1

Cuthbertson, K., Hayley, S., Motson, N., and Nitzsche, D. (2016). What does rebalancing really achieve? *International Journal of Finance and Economics*, 21(3), 224-240. https://doi.org/10.1002/ijfe.1545

Erb, C. B. and Harvey, C. R. (2006). The strategic and tactical value of commodity futures. *Financial Analysts Journal*, 62(2), 69-97. https://doi.org/10.2469/faj.v62.n2.4084

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25(2), 383-417. https://doi.org/10.1111/j.1540-6261.1970.tb00518.x

Fama, E. F. (1991). Efficient capital markets: II. *Journal of Finance*, 46(5), 1575-1617. https://doi.org/10.1111/j.1540-6261.1991.tb04636.x

Feld, A. (1999). High exposure to low-basis stock: Too much of a good thing? *CPA Journal*, 69(11), 60-65.

Hasbrouck, J. (2009). Trading costs and returns for U.S. equities: Estimating effective costs from daily data. *Journal of Finance*, 64(3), 1445-1477. https://doi.org/10.1111/j.1540-6261.2009.01469.x

Malkiel, B. G. (2003). The efficient market hypothesis and its critics. *Journal of Economic Perspectives*, 17(1), 59-82. https://doi.org/10.1257/089533003321164958

Smith, M. H., and Smith, G. (2008). Harvesting capital gains and losses. *Financial Services Review*, 17(4), 309-321.

Statman, M. (1987). How many stocks make a diversified portfolio? *Journal of Financial and Quantitative Analysis*, 22(3), 353-363. https://doi.org/10.2307/2330969

Stein, D. M., Siegel, A. F., Narasimhan, P., and Appeadu, C. E. (2000). Diversification in the presence of taxes. *Journal of Portfolio Management*, 27(1), 61-71.

Whitworth, J. (2018). Improving long-term portfolio risk and return by using appreciated stocks for charitable donations. *Financial Services Review*, 27(3), 257-277.

Willenbrock, S. (2011). Diversification return, portfolio rebalancing, and the commodity return puzzle. *Financial Analysts Journal*, 67(4), 42-49. https://doi.org/10.2469/faj.v67.n4.1