

Tax-loss harvesting: Why a personalized approach is important

- The value of tax-loss harvesting (TLH) varies significantly across investor characteristics, investor behavior, and market environment. Each of these three categories drives roughly one-third of the variation in value. As a result, tax-loss harvesting is not without risk, particularly for investors with lower tax rates and less in capital gains to offset. These investors can experience modest benefits or, in some cases, negative effects.
- We consider several net-worth-based profiles informed by the Survey of Consumer Finances. Every profile analyzed showed improvement in TLH outcomes when simulating an optimal TLH program, relative to suboptimal behavior. The single most important behavior driving the value of tax-loss harvesting is the reinvestment of tax savings in the portfolio.
- Grounded by the net-worth-based investor profiles, we find “optimal” TLH alpha—the excess annual after-tax return achieved by conducting TLH—to range between 0.47% and 1.27%.¹

Authors



Thomas Paradise,
CFA



Kevin Khang,
Ph.D.



Joel M. Dickson,
Ph.D.

¹ These numbers reflect a 15-year projected excess return from TLH as measured by the increase in internal rate of return. To understand the potential impact of TLH for an investor, the value reported must be scaled by the size of taxable equity assets relative to the size of the entire portfolio. That is, if an investor's TLH alpha is 1%, but the applicable assets (taxable equity) reflect only 25% of the portfolio, the expected impact to the portfolio would be only 0.25%.

Introduction

Tax-loss harvesting is a widely used strategy in the field of personal wealth management, often lauded for its potential to improve an investor's after-tax return. When executed within a well-designed investment plan, TLH can improve the efficiency of a portfolio by deferring, and in some cases eliminating, tax payments, which keeps more capital invested and compounding. For this reason, TLH is a strategy that closely aligns with the Vanguard principle of minimizing cost (Vanguard, 2023), which is central to enhancing investor returns.

The value provided to investors, however, can vary significantly. It is driven by a multitude of factors that can be grouped into three categories: characteristics specific to the investor, things that the investor can control or influence, and market behavior outside the investor's control. This analysis aims to attribute the relative importance of factors on TLH outcomes and discuss their importance through both qualitative and quantitative lenses. We use this understanding of how TLH value varies and an analysis of investor profiles to answer two questions: What is a reasonable expectation of value from TLH personalized to the individual? And what can investors do to maximize their potential benefit?

Building on previous research (Khang, Paradise, and Dickson, 2021), we use historical simulation combined with machine-learning techniques to attribute TLH value to core drivers. To account for the inherent uncertainty and variability in market influences, we consider a probabilistic approach when evaluating the outcomes of TLH, presenting results as a distribution of outcomes. This approach diverges from the traditional method of presenting a single-point estimate of performance (for example, see Berkin and Ye, 2003, and Chaudhuri, Burnham, and Lo, 2020).

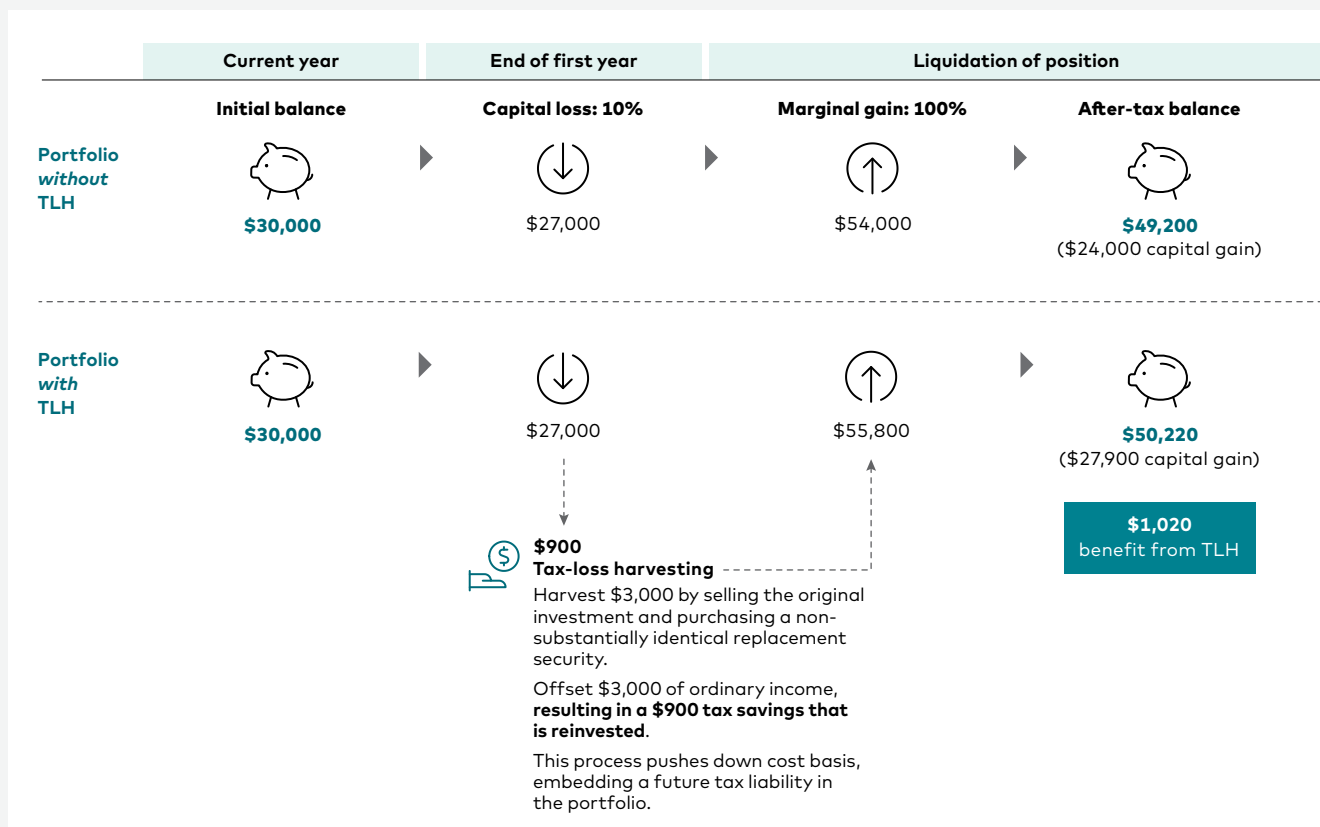
We further extend our previous work by examining how the investment horizon influences the expected value derived from TLH, providing insights into the temporal aspect of the strategy's efficacy. This temporal aspect is particularly relevant for tailoring advice to investors with varying time horizons and cash-flow constraints.

How TLH works

Tax-loss harvesting operates on the principle of converting investment losses into tax savings. Securities held in a taxable account can be sold—or “harvested”—when their value falls below cost basis (the value when they were purchased); the investor thus realizes a capital loss. These capital losses are used to offset an equal amount of capital gains in other parts of an investor’s portfolio and/or up to \$3,000 of ordinary income. If the losses realized exceed the gains realized and \$3,000 of ordinary income in a given calendar year, the extra losses are carried forward and are available to offset capital gains and income in future years.

The following example (Figure 1) illustrates how TLH can add value. When the portfolio experiences a loss, the position can be sold, and the captured loss (\$3,000) can be used to offset income, creating a tax savings of \$900. When this \$900 is reinvested in the portfolio, the portfolio grows more than it would have without the reinvestment of tax savings. Ultimately, the TLH portfolio can also generate a larger tax bill because of its lower cost basis; however, this will represent a net benefit if the market appreciates after the reinvestment of tax savings and/or the tax rate paid by the investor is lower at liquidation than at the time of harvest.

FIGURE 1
TLH adds value through tax-rate substitution and increased market exposure



Notes: This hypothetical illustration does not represent the return on any particular investment and the rate is not guaranteed. This figure assumes a 30% tax rate when losses are harvested and a 20% tax rate when the position is liquidated.

Source: Vanguard.

What factors drive TLH value?

The value of TLH hinges on three foundational components: the ability to generate losses, the ability to convert those losses into tax savings efficiently, and the ability to reinvest tax savings in the stock market—which historically has appreciated over the long run. This framing is important because an investor needs to achieve all three components to take full advantage of tax-loss harvesting. It is primarily for this reason that the benefits of TLH are not equal for all investors.

For instance, if an investor's portfolio is not exposed to volatility, the investor cannot harvest losses to offset taxes. If an investor can generate losses, but does not have sufficient capital gains and income to offset, or does not have a high-enough tax rate, the investor cannot convert

those losses to tax savings in the current tax year, which may reduce the effectiveness of TLH. Finally, even if an investor can harvest losses and convert them efficiently into tax savings, that investor needs to have the discipline to invest those tax savings in the market to fully reap the rewards of TLH.

In **Figure 2**, we break down the three foundational components into factors and explain how each influences the value of TLH. To derive the relative importance of each of the factors, we used a historical TLH simulation model paired with machine-learning techniques to quantify how each factor contributed to the variance of TLH value for periods going back to 1982. More details on the simulation and attribution techniques can be found in the **Appendix**.

FIGURE 2

Robust TLH benefit requires all three drivers: Loss generation, conversion to tax savings, and increased market exposure

Drivers	Factors		Relative importance
Loss generation 31%	Volatility	Security volatility, especially shortly after investment, drives prices below their cost basis, enabling loss harvesting.	17%
	Harvest frequency	Monitoring the portfolio regularly for harvests enables a greater capture of ephemeral losses (Khang, Cummings, and Paradise, 2022).	7%
	Recurring investments	New investments represent new tax-loss-harvesting opportunities because they create new lots more likely to fall below their cost basis in an appreciating market.	2%
	Portfolio granularity (e.g., direct indexing)	In addition to broad market volatility, cross-sectional volatility between securities, unlocked by owning small portfolio building blocks (for example, direct indexing), increases loss-harvesting opportunities.	3%
	Time horizon	As markets appreciate and more losses are harvested from a portfolio, the relative cost basis of the portfolio is lower, making future harvest opportunities rarer.	2%
Conversion to tax savings 32%	Current and future taxes	The higher the current tax rate, the more tax savings are generated by offsetting income. The lower the future tax rate, the lower the embedded tax liability created by TLH. Additionally, the degree to which the taxable portfolio is eventually liquidated is proportional to the amount in taxes generated by TLH because of the step-up in cost basis provision. The basis for inherited assets is their appraised value at the date of death. See IRS Publication 551 for additional information.	27%
	Loss offsetting income	Capital gains are required to turn loss harvests into tax savings (Khang, Paradise, and Dickson, 2021). Without capital gains, taxable income can be reduced by up to just \$3,000 of ordinary income.	5%
Market exposure 37%	Reinvesting tax savings	The more an investor reinvests tax savings, the more efficiently the investor leverages deferred taxes to increase market exposure. In addition, the investor is creating more opportunities for future harvests.	25%
	Market return	The compounding of reinvested tax savings is proportional to the market return after reinvestment. On the other hand, steadily appreciating markets reduce opportunities for tax-loss harvesting.	12%

Notes: Relative importance was determined by SHAP (SHapley Additive exPlanations) analysis. More details can be found in the Appendix. Data for the SHAP analysis was generated systematically using approximately 10,000 simulations over historical market returns from January 1982 through March 2023. The following variables were randomized using uniform distributions to create the simulations: harvest tax rate, liquidation tax rate, loss offsetting income, portfolio granularity, harvest frequency, quarterly investment amount, level of tax savings reinvestment, start date, and time horizon. The relative importance of the variables noted in Figure 2 is generalized across a wide swath of hypothetical profiles. Because of nonlinear relationships between TLH alpha and these variables, the importance of some drivers will vary by investor. For instance, when we consider profiles of high-net-worth investors, the importance of direct indexing was 9%. This relationship is discussed more later.

Sources: Vanguard calculations, using data from Axioma.


It's about the amount of taxes you save, not the amount of losses you harvest

For each investor, either the ability to generate losses or the ability to convert losses into tax savings will determine whether the investor can further benefit from TLH. Investors will want to take different courses of action to expand one capacity or the other. In some cases, investors will not have sufficient losses to offset all their capital gains. This is common for certain (ultra) high-net-worth investors who tend to generate capital gains income from various asset holdings. In these cases, they will want to generate more losses by increasing portfolio granularity with direct indexing or increasing harvesting intensity.

However, in alternative scenarios, investors may have more than sufficient capital losses to offset their available capital gains, or they may expect their taxes to rise in the future. This is common for primarily passive investing accumulators whose products are tax-efficient—they do not generate significant gain distributions or regularly draw down their portfolios to create capital gains. In either situation—insignificant gain distributions or lack of regular drawdowns—these investors may not be suitable candidates for interventions aimed at increasing loss harvests, such as direct indexing. If these investors already have sufficient losses to offset all their gains and \$3,000 of income annually, increasing harvests may yield no benefit in the current period, while introducing additional transaction costs and tracking error into the portfolio. If the investors expect their taxes to be higher in the future, any harvests may only subtract value. It is important for investors to consider how their situations may change over time. Even if they do not need additional harvests right now, those losses can be carried forward to future years to offset future ordinary income and capital gains.

The personalized nature of TLH benefit is best illustrated in the context of what is specific to the investor, what an investor can influence, and what is outside an investor's control. In this framework, using the relative importance of the drivers in Figure 2, we find that each component accounts for roughly a third of the outcome (Figure 3). The biggest component drivers are whether investors are disciplined about reinvesting the tax savings yielded from the TLH process and what their current and future tax rates are.

FIGURE 3
The benefit of TLH depends on three broad, roughly equal components



Notes: Relative importance was determined from SHAP analysis. More details can be found in the Appendix. Data for the SHAP analysis were generated systematically using approximately 10,000 simulations over historical market returns from January 1982 through March 2023. The following variables were randomized using uniform distributions to create the simulations: harvest tax rate, liquidation tax rate, loss offsetting income, portfolio granularity, harvest frequency, quarterly investment amount, level of tax savings reinvestment, start date, and time horizon. Features specific to the investor include time horizon, current and future tax rates, and loss offsetting income. Features an investor can influence include recurring investments, portfolio granularity, harvest frequency, and reinvesting tax savings. Features an investor can't influence include volatility and market return. Percentages do not add up to 100% because of rounding.

Sources: Vanguard calculations, using data from Axioma.

Risks and other key considerations associated with TLH

It is important to note that TLH is not a free lunch. There are risks for any investor engaging in TLH and it can even subtract value in certain circumstances. We review these risks along with other issues:

Wash sales

A critical limitation on the application of TLH is the wash sale rule—a provision designed to prevent investors from abusing TLH by realizing losses while maintaining exposure to the harvested securities or those that would behave nearly identically. The rule prohibits investors from claiming a loss on a sale if they purchase a “substantially identical” security within 30 days before or after the sale.²


To mitigate this risk, investors must carefully select replacement securities that maintain a desirable but distinct exposure to avoid running afoul of the rule. Notably, this is mainly a risk for investors who conduct TLH in a portfolio with commingled funds, such as mutual funds and ETFs, which may have several replacements that behave nearly identically.

Ossification

When securities are harvested and the proceeds are reinvested, the cost basis of the portfolio is pushed down. This, in combination with the upward historical trend of markets, makes those assets from the reinvested proceeds harder to harvest in the future, which can manifest in dwindling TLH alpha over the investment horizon (Figure 4). This hardening of the portfolio against harvest opportunities is commonly known as “ossification.” Investors who find this property undesirable may broaden their harvesting opportunities by investing dividends, tax savings, and external cash flows in the portfolio to create new tax lots at higher cost bases (Khang, Paradise, and Cummings, 2023). These new tax lots are more likely to be harvestable with normal market volatility.

It is also important to note that even portfolios with very low cost bases can be “de-ossified” during extreme market downturns such as those in 2008 and 2020. The need to prepare for the possibility of ossification is greatest among investors well-suited to TLH—those who have regular capital gains to offset and who expect a high TLH benefit.

FIGURE 4
TLH value can diminish over time without diligent cost basis renewal



Notes: The TLH simulations use historical market returns from January 1982 through March 2023. The distribution of projected outcomes is determined using rolling time periods of the same length. This chart uses profile assumptions consistent with Profile 3 described in Figure 7, on page 9. Tax savings are calculated assuming that two-thirds of offset income is subject to long-term capital gains tax rates and one-third is short-term gains subject to ordinary income tax rates. The simulations include scanning for harvests daily, reinvesting all tax savings in the portfolio, making quarterly contributions equal to 10% of the initially invested principal, and harvesting in a direct-indexed portfolio of 400 securities. TLH alpha numbers are annualized over the simulation period.

Sources: Vanguard calculations, using data from Axioma.

² See IRS Revenue Ruling 2008–5 for additional information.

Capital gains availability

Investors need sufficient access to capital gains to convert loss harvests into tax savings, and the aggressiveness of their tax-loss-harvesting program should be calibrated based on capital gains availability (Khang, Paradise, and Dickson, 2021). This can represent a risk because it can be difficult to predict the amount and timing of capital gains in the future. One way to mitigate this risk is to consider what future events may cause significant gain realization and consider them as part of a holistic financial plan. For example, selling a sizable amount of investments to make a large purchase such as a house or selling a personal business represent foreseeable tax-planning events that could be considered in a TLH plan.

FIGURE 5
2024 0% tax rates

Type	Single	Married filing jointly
Long-term capital gains	Up to \$61,625*	Up to \$123,250*
Ordinary income/short-term capital gains	Up to \$14,600*	Up to \$29,200*

* Assumes \$14,600 standard deduction for single filers and \$29,200 for married filers, filing jointly.

Note: If an asset is held for more than one year before it is sold, the capital gain or loss is considered long-term. See IRS topic number 409 for more information.

Sources: Vanguard calculations, using data from the IRS.


Tax rate changes

As discussed earlier, deferring taxes to a point when tax rates are higher can subtract value. The easiest way to mitigate this is to not engage in tax-loss harvesting when subject to “0% tax rates” (**Figure 5**)—that is, the capital gains or the income are not high enough to trigger a tax. However, many investors can find it challenging to predict their future tax rates if they expect their circumstances or earning power to change over time. In addition, the tax code can always change, potentially leading to higher future tax rates or the creation of other provisions to limit the efficacy of TLH.

Market environment

Ultimately, TLH value will be amplified or dampened by the market return over the investment horizon. One way to think of the TLH process is in terms of leverage. By harvesting a loss, investors are borrowing money from their future selves in the form of tax deferral. As with leverage, the benefit they receive from this “borrowed” money is directly proportional to the market return they can achieve with the invested tax savings. Typically, in rising markets, this translates to a benefit; however, investors who exit the market after a sharp downturn, concluding TLH-based investing, can be harmed by tax-loss harvesting because they increased their exposure to a negative market. **Figure 6** illustrates the variability in TLH outcomes purely as a function of market timing.

FIGURE 6
TLH benefits vary with market performance



Notes: The TLH simulations use 10-year rolling periods over historical market returns from January 1982 through March 2023. This chart uses profile assumptions consistent with Profile 3 described in Figure 7. Tax savings are calculated assuming that two-thirds of offset income is subject to long-term capital gains tax rates and one-third is subject to ordinary income tax rates. The simulations include scanning for harvests daily, reinvesting all tax savings in the portfolio, making quarterly contributions equal to 10% of the initially invested principal, and harvesting in a direct-indexed portfolio of 400 securities. TLH alpha numbers are annualized over the simulation period.

Sources: Vanguard calculations, using data from Axioma.

Past performance is no guarantee of future results.

Profile analysis

Armed with an understanding of the driving features of TLH and their relative importance, we can illustrate expected ranges of TLH value through varying key investor attributes and behaviors. We anchor this analysis in the Survey of Consumer Finances (SCF) to give a realistic expectation of potential value based on client net worth and associated demographic information.

The SCF is a longitudinal study conducted by the Federal Reserve. It provides comprehensive data on the financial conditions of U.S. households including income, net worth, and asset ownership. We use it to create realistic representations of different investor types across net-worth groups (Figure 7).

FIGURE 7
Four profiles representing net-worth groups

Type	Profile 1	Profile 2	Profile 3	Profile 4
Net-worth group (percentiles)	75th to 90th	90th to 95th	95th to 98th	Top 2%
Ordinary income tax rate	22.0%	24.0%	41.3%	48.3%
Long-term capital gains tax rate	15.0%	15.0%	24.3%	31.3%
Offsettable income	2%	4%	6%	9%

Notes: Profiles 1 and 2 are assumed to have no additional state income taxes. Profile 3 is assumed to be subject to a 32% federal marginal income tax rate, a 15% federal marginal long-term capital gains tax rate, and a 9.3% California marginal capital gains tax rate. Profile 4 is assumed to be subject to a 37% federal marginal income tax rate, a 20% federal marginal long-term capital gains tax rate, and an 11.3% California marginal capital gains tax rate.


Sources: Vanguard calculations, using data from the Survey of Consumer Finances.

By simulating TLH alpha for these profiles, we can identify reasonable expectations for investors in these net-worth groups given a variety of historical economic scenarios. In addition, we can show how optimal behavior will vary outcomes.

Figure 8 shows the expected outcomes for TLH over time. The left side shows how each net-worth group is projected to fare when engaging in suboptimal TLH behavior. This includes quarterly harvesting, reinvesting 50% of tax savings, no ongoing cash flows, and investing in a commingled fund. The optimal behavior considered includes harvesting daily, reinvesting all tax savings, making quarterly contributions, and harvesting in a direct-indexed portfolio—many of which are covered in Khang, Cummings, and Paradise (2022).

As shown, higher-net-worth investors are more likely to benefit from TLH. The median 15-year projected annualized TLH alpha across these net-worth profiles was 0.47%, 0.65%, 1.00%, and 1.27%, respectively, in the context of their taxable equity portfolio when simulating optimal behavior. This compares with 0.04%, 0.05%, 0.10%, and 0.13% when simulating suboptimal behavior. We also see ossification, the pattern described earlier in the paper, at work—TLH alpha diminishes over time. It is most notable for high-net-worth investors well-suited to TLH. Finally, we can see that any net-worth group can have a significant impact on its expected TLH alpha by engaging in optimal behaviors such as daily harvest screening, reinvesting tax savings, making ongoing contributions to the portfolio, and using a more granular portfolio when suitable.

FIGURE 8
Range of TLH alpha by net-worth group and investor behavior



Notes: The TLH simulations use historical market returns from January 1982 through March 2023. The distribution of projected outcomes is determined using rolling time periods of the same length. This chart uses profile assumptions described in Figure 7. Tax savings are calculated assuming that two-thirds of offset income is subject to long-term capital gains tax rates and one-third is subject to ordinary income tax rates. The left side of the chart represents suboptimal behavior including scanning for harvests quarterly, reinvesting 50% of tax savings in the portfolio, making no quarterly contributions, and harvesting in a single commingled fund. The right side of the chart represents optimal behavior including scanning for harvests daily, reinvesting all tax savings in the portfolio, making quarterly contributions equal to 10% of the initially invested principal, and harvesting in a direct-indexed portfolio of 400 securities. TLH alpha numbers are annualized over the simulation period.

Sources: Vanguard calculations, using data from Axioma.

Case study

To further illustrate the potential value and variance of TLH alpha resulting from tax-loss harvesting, we consider two hypothetical cases.

Mass-affluent investor


The first case features an affluent investor (75th–90th-percentile net worth). This investor primarily invests in passive, tax-efficient, commingled ETFs. Approximately 25% of this investor's portfolio is invested in taxable equity assets. To illustrate how market conditions and luck will influence an investor's outcomes, we use rolling historical periods to create a distribution of projected outcomes (Figure 9). We show two sets of projections. The first indicates the range of expected TLH alpha on taxable equity assets. The second indicates the range of expected TLH alpha when scaled relative to the entire portfolio.

Note that the range of benefit that any individual investor will receive can vary dramatically as evidenced by the range of values between 10th-percentile outcomes and 90th-percentile outcomes. Even if an investor is well-suited to

TLH and acts to maximize the potential for success, market volatility, return, and luck will drive a significant portion of the realized TLH value, particularly in the near term. The expected 15-year annualized TLH alpha for this investor is 47 basis points. (A basis point is one-hundredth of a percentage point.) This number, however, applies only to taxable equity assets. When TLH alpha is considered in the context of the investor's entire portfolio—taking into account the proportion of taxable equity assets in the portfolio—this value drops to 12 basis points.

More than anything else, this case study illustrates the importance of looking at TLH benefit from the perspective of the entire portfolio and setting reasonable expectations when comparing investment options. For investors ideally suited to TLH, the commonly discussed TLH benefit on the order of 100 basis points may be relevant for the entire portfolio.³ But there are many other investors for whom TLH benefit in the context of the entire portfolio may be materially smaller.

FIGURE 9
TLH alpha should be considered in the context of the entire portfolio



Notes: The TLH simulations use historical market returns from January 1982 through March 2023. The distribution of projected outcomes is determined using rolling time periods of the same length. This chart uses profile assumptions consistent with Profile 1 described in Figure 7. Tax savings are calculated assuming that two-thirds of offset income is subject to long-term capital gains tax rates and one-third is subject to ordinary income tax rates. The simulations include scanning for harvests daily, reinvesting all tax savings in the portfolio, making quarterly contributions equal to 10% of the initially invested principal, and harvesting in a direct-indexed portfolio of 400 securities. TLH alpha numbers are annualized over the simulation period.

Sources: Vanguard calculations, using data from Axioma.

³ Depending on the investor profile, there may be additional portfolio construction implications for the entire portfolio (Khang, Cummings, Paradise, and O'Connor, 2022).


High-net-worth investor with expected sale of business

A common use case of tax-loss harvesting is to mitigate the realized gains from a significant transaction such as the sale of real estate or a business (Figure 10). This case features a high-net-worth business owner (95th–98th-percentile net worth). The business owner primarily invests in individual securities to maximize the losses achievable from TLH and to minimize the investment portfolio’s exposure to the industry in which the owner does business. The business owner does not regularly realize capital gains but

expects to sell the business in five years, realizing gains representing 100% of the owner’s taxable equity assets.

By engaging in TLH, this business owner is projected to improve the portfolio return by approximately 50 basis points following the sale of the business. This scenario also illustrates that even investors who may not benefit in the short term from TLH may still wish to harvest losses if they are likely to be suitable candidates in the future.

FIGURE 10
TLH can mitigate the tax impact of a sale with large capital gains implications



Notes: The TLH simulations use historical market returns from January 1982 through March 2023. The distribution of projected outcomes is determined using rolling time periods of the same length. This chart uses profile assumptions consistent with Profile 3 described in Figure 7. Prior to five years, no additional offsettable income is assumed. At Year 5, loss offsetting income equal to 100% of the account is assumed. After Year 5, 6% loss offsetting income is assumed, consistent with Profile 3. Tax savings are calculated assuming that two-thirds of offset income is subject to long-term capital gains tax rates and one-third is subject to ordinary income tax rates. The simulations include scanning for harvests daily, reinvesting all tax savings in the portfolio, making quarterly contributions equal to 10% of the initially invested principal, and harvesting in a direct-indexed portfolio of 400 securities. TLH alpha numbers are annualized over the simulation period.

Sources: Vanguard calculations, using data from Axioma.

Conclusion

This article explored personalized expectations for tax-loss harvesting and the behaviors investors can emulate to increase their odds of success. We showed that TLH alpha can range widely, depending on investor characteristics, behavior, and return environment. We quantified the relative importance of these main determinants in driving the TLH alpha, finding that the investor has influence over the greatest amount of alpha variation (37%). Another 34% is driven by characteristics specific to the investor, and the final 28% is based on market volatility and return—factors largely outside the investor's control. (These numbers add up to only 99% because of rounding.)

Given the significant dispersion around the potential TLH outcomes and the importance of factors within their control, what should individual investors do? Our analysis of investor profiles suggests that investors should engage in TLH with reasonable expectations of what can be achieved. TLH is not without risks, particularly for investors subject to low taxes in the current period. The most important thing that any investor can do to maximize the potential for TLH success is to reinvest tax savings in the portfolio.

References

- Aladangady, Aditya, Jesse Bricker, Andrew C. Chang, Sarena Goodman, Jacob Krimmel, Kevin B. Moore, Sarah Reber, Alice Henriques Volz, and Richard Windle, 2023. *Changes in U.S. Family Finances From 2019 to 2022: Evidence From the Survey of Consumer Finances*. Washington, D.C.: Board of Governors of the Federal Reserve System.
- Berkin, Andrew L., and Jia Ye, 2003. Tax Management, Loss Harvesting, and FIFO Accounting. *Financial Analysts Journal* 59(4): 91–102.
- Chaudhuri, Shomesh E., Terence C. Burnham, and Andrew W. Lo, 2020. An Empirical Evaluation of Tax-Loss Harvesting Alpha. *Financial Analysts Journal* 76(3): 99–108.
- Khang, Kevin, Thomas Paradise, and Joel Dickson, 2021. Tax-Loss Harvesting: An Individual Investor's Perspective. *Financial Analysts Journal* 77(4): 128–150.
- Khang, Kevin, A. Cummings, and Thomas H. Paradise, 2022. Optimal Tax-Loss-Harvesting Implementation. *The Journal of Beta Investment Strategies, Direct Indexing* 2022, 13(2): 75–89.
- Khang, Kevin, Alan Cummings, Thomas Paradise, and Brennan O'Connor, 2022. *Personalized Indexing: A Portfolio Construction Plan*. Valley Forge, Pa.: The Vanguard Group.
- Khang, Kevin, Thomas Paradise, and Alan Cummings, 2023. Expected Loss Harvest From Tax-Loss Harvesting With Direct Indexing. *The Journal of Beta Investment Strategies, Direct Indexing* 2023, 14(3): 42–56.
- Vanguard Group, 2023. *Vanguard's Principles for Investing Success*. Valley Forge, Pa.: The Vanguard Group.

Appendix

Data and tax-loss-harvesting algorithm

Constituents of our U.S. equity capitalization-weighted index come from the top 400 securities by market capitalization in the Axioma US4 risk model, which we refer to as "Axioma 400." This index reconstitutes on the first day of each calendar year, using market capitalizations as of the last day of the prior calendar year. Securities that would be subject to mergers and acquisitions in the upcoming calendar year are excluded from consideration for Axioma 400. Daily returns in the Axioma 400 universe range from the beginning of 1982 to the end of March 2023.

To construct different levels of aggregated securities, securities are grouped randomly into aggregate assets of equivalent counts of securities. Each aggregate security maintains the summed market capitalization of its components so that the market return of the entire portfolio is maintained.

Our decision to use this synthetic index instead of the Standard & Poor's 500 Index or the Russell 1000 Index is primarily driven by the ease of accessibility to daily information on returns and dividends at the individual-security level. In terms of aggregate index-level returns, Axioma 400 behaves very similarly to the S&P 500 Index, the Russell 1000 Index, and the Russell 3000 Index, with an R-squared of 95% to 96%, β of 0.99 to 1.01, and α of -0.01 to 0.01 when we regress Axioma 400 returns on these other indexes.

Our TLH algorithm maintains individual tax lot holdings of all individual-security positions daily. We consider every security to have a single perfect replacement. We can sell one security in a pair and buy the other. That said, our algorithm does not trade a security if either its sale or the purchase of the replacement would result in a wash sale. In cases where neither the original security nor the single perfect replacement is available for purchase because of wash sale considerations, the weights of other available securities are proportionally increased. This approach isolates the availability of losses from the effects of tracking error in navigating the wash sale rule. The effect of this treatment may not be neutral, so this guides our interpretation of the loss-harvest results to be indicative of the loss-harvesting opportunity existing throughout the 40-year history we examine.

The TLH portfolio receives an initial cash contribution at the beginning of the period, and at the beginning of subsequent quarters based on the cash contribution scenario. During the run, all subsequent cash infusions, dividends, and proceeds from security sales are kept in cash until either the beginning of the next quarter or the next loss harvest, whichever comes first, when all available cash is invested toward the target allocation of the portfolio.

Loss harvesting is initiated if a given tax lot is in breach of a loss threshold of 5% relative to the cost basis. We keep track of the portfolio's cumulative loss harvest, net of any capital gains realized because of index reconstitution.

To quantify the value of tax-loss harvesting and its primary drivers, we use an XGBoost (eXtreme Gradient Boosting) regression, a machine-learning technique that uses an ensemble of decision trees, optimized through gradient boosting, to predict the input-output relationship. It is well-suited to our purposes because of its ability to handle nonlinear and complex interactions between variables relevant to the value of TLH.

By training an XGBoost model on our data, we can understand the relative importance of driving factors based on how much they contribute to the model's predictive accuracy, giving us a clear indication of which factors are most influential in determining the value of TLH. We can also begin

to understand the direction and magnitude of each feature's impact on the value of TLH both independently and in connection with other features through SHAP (SHapley Additive exPlanations) analysis and partial dependence plots. Ultimately, we can use the trained model to predict personalized value of TLH for individual investors in different scenarios more efficiently than simulating their performance independently.

We derive the relative importance of our drivers using SHAP analysis. SHAP values provide a prediction for the marginal contribution of each feature on the TLH alpha⁴ outcome using a fair distribution of values for the other considered features. This is distinct from traditional feature importance, which accounts for the average effects of other features by measuring each feature's impact on model error. In this regard, SHAP value better considers the interaction between features.

⁴ In this analysis, TLH alpha is measured using a time-weighted return (TWR). This measure is useful for understanding how factors influence the efficacy of TLH because it removes the influence of timing of cash flows, making it easier to understand what is driven by factors versus the idiosyncrasies of market-timing. It is important to note that TWR does not effectively measure the return experienced by the investor. An alternative metric that more accurately captures the investor return over time is differential internal rate of return (IRR), which is used in the rest of the analyses in this paper. IRR, however, is sensitive to the cash-flow timing idiosyncrasies discussed above.

Connect with Vanguard®
vanguard.com

All investing is subject to risk, including possible loss of principal.

Tax-loss harvesting involves certain risks, including, among others, the risk that the new investment could have higher costs than the original investment and could introduce portfolio tracking error into your accounts. There may also be unintended tax implications. We recommend that you carefully review the terms of the consent and consult a tax advisor before taking action.

CFA® is a registered trademark owned by CFA Institute.

Vanguard®

© 2024 The Vanguard Group, Inc.
All rights reserved.

ISGTLH 072024