First Trust

Analysis of Strategies for Managing Direct Indexing Accounts

Abstract:

This paper analyzes various strategies in managing direct indexing (DI) accounts with an emphasis on tax loss harvesting (TLH) and tracking error (TE) concerns. The first topic analyzes the novel approach of using a tradeoff coefficient to combine the competing objectives of maximizing TLH and minimizing TE. The second section explores the general question, "How much TLH can be expected over time?" As an answer to this question, many authors point to the phenomenon of ossification – how a DI account with no cash inflows captures tax losses less efficiently over time. However, the result of our study uncovers far more nuanced possibilities than those obtained by other authors who use Markov chain simulations to analyze ossification. Finally, this paper addresses one strategy to mitigate the ossification in a DI account, namely a pattern of regular cash infusions.

Important Disclosures

This white paper analyzes the implementation and performance of a Direct Indexing investment strategy ("DI Strategy"). This white paper is for information purposes only and is not intended as financial, legal or tax advice. First Trust does not provide legal, tax or accounting advice and does not assume any liability for the financial, legal or tax consequences of any securities transactions. DI strategy investors should consult with their personal financial, tax or legal advisor regarding the tax consequences of investing within a direct indexing investment strategy, based on their particular circumstances. DI strategy investors and their personal tax advisors are responsible for how the transactions conducted in an account are reported to the IRS or any other taxing authority on the investor's personal tax returns. First Trust assumes no responsibility for tax consequences to any investor of any transaction.

The effectiveness of the tax-loss harvesting strategy to reduce the tax liability will depend on the DI Strategy investor's entire tax and investment profile, including purchases and sales in accounts not invested in a DI Strategy, the type of investments (e.g., taxable or nontaxable) or holding period (e.g., short-term or long-term). The performance of new securities purchased through tax-loss harvesting may be better or worse than the performance of securities sold for tax-loss harvesting purposes.

This material may not be reproduced, in whole or in part, without the written consent of First Trust Advisors L.P. ("First Trust"). First Trust and its affiliates are not responsible for its use by other parties. The information presented in this white paper is intended solely to analyze DI strategies. It is for illustrative purposes only and not indicative of any actual portfolio. We believe the information provided is reliable but do not warrant its accuracy or completeness. **This material is not intended as an offer or solicitation to invest in a DI Strategy or buy or sell a specific security.** Any mention of specific securities is intended to help demonstrate the objectives of a DI Strategy and should not be considered a recommendation or solicitation to purchase or sell securities**. It should not be assumed that any investment in a DI Strategy was, or will prove to be, profitable, or that the investment decisions made in the future will be profitable or current performance of a DI Strategy will persist. Past performance is not indicative of future results.** This material does not take into account any investor's particular investment objectives, strategies, tax status, or investment horizon. A DI Strategy described herein may not be suitable for all investors. Investing entails risks, and there can be no assurance that a DI Strategy will achieve profits or avoid incurring losses.

First Trust makes no representation or endorsement concerning the accuracy of information received from any third party. The opinions and estimates expressed in this report are those of the authors and are current only through the date stated at the top of this page. These views are subject to change at any time based on market trends or other conditions, and First Trust disclaims any responsibility to update such views. **These views may not be relied on as investment advice and, because investment decisions are based on many factors, may not be relied on as an indication of trading intent on behalf of any DI strategy.**

Contents

Introduction

Background on Direct Indexing

For nearly a century, mutual funds have offered investors the potential risk-mitigation benefits of diversifying a portfolio while removing the burden of managing the constituent securities. By commingling a basket of underlying securities into a single tradable vehicle, exchange-traded funds (ETFs), first created in the 1990s, added the potential advantages of liquidity and improved tax efficiency of investing in funds. The first DI accounts also date to the early 1990s but recent technological improvements in data management and trading infrastructure, which lowered trading costs, established an environment that made DI possible for more investors.

Unlike mutual funds and ETFs, in a DI paradigm, the investor directly owns the underlying securities in a separately managed account (SMA). Using the recently developed tools, the DI account manager can easily make dozens or even hundreds of trades at once without incurring significant trading costs. This framework gives the investor far more flexibility in crafting a portfolio. Among the many potential benefits this affords, four large categories stand out.

1. **Tracking an index**: By definition, DI accounts track a specific reference index. By doing so, DI accounts retain the diversification benefits of mutual funds and ETFs. Industry professionals sometimes say that direct indexing “unwraps the ETF.” An account manager could attempt to perfectly mirror a reference index, mimicking how it rebalances from time to time. To *track an index* refers to a more technical process by which the account manager regularly rebalances the portfolio to be close to the reference index, where a so-called *tracking error* (TE) defines the measure of closeness of the portfolio's holdings to the constituents of the reference index. We discuss some technical aspects of this in the next subsection.
2. **Customization**: Investors may wish to craft a portfolio along certain individual preferences. Some might want to tilt their portfolio towards companies with desired factors or fundamentals. Other investors may care about the values-based qualities of companies in their portfolio. Indeed, clients increasingly want to invest in companies that have a positive impact on environment, society, and governance according to their personal values. Some investors may already have considerable exposure to a single security and may wish their DI account to exclude that security. For example, if an investor began their career at Hewlett-Packard and acquired shares through a stock options program, that person might choose not to have Hewlett-Packard in their DI portfolio or may try to minimize TE around the appreciated position. Regardless of the reason, DI affords customization options.
3. **Tax loss harvesting (TLH)**: Arguably the most quantifiable benefit of a DI account is the ability to manage after-tax outcomes. The tax code does not allow individual investors to claim capital losses realized within mutual funds or ETFs, and an investor can only claim capital losses when the mutual fund or ETF decreases in value below the initial invested amount. Not so when the investor directly owns the securities: he or she can claim losses from a single security or tax lot that drops below its initial cost basis if that tax lot or security is sold. Following a certain rebalancing frequency or perhaps strategically triggered by certain events, a DI manager can rebalance a portfolio according to the following steps: (a) sell securities that decline in value, thereby realizing capital losses; (b) immediately reinvest the proceeds in a manner that keeps the modified portfolio within the tracking error of the reference index; (c) utilize the capital losses to offset capital gains realized elsewhere in the investor’s portfolio. This process is called *tax loss harvesting*. Furthermore, the tax code allows the investor to carry losses forward into future years.
4. **Portfolio Transitions**: Another tax-related benefit of DI arises during portfolio transitions. When an investor transitions a portfolio from an aggressive to a conservative strategy, or from one values-based investing tilt to another, or away from an undesired concentrated position, DI can help achieve a more tax-efficient transition.

The flexibility of a DI account inevitably carries a flip side: more responsibility on the part of the account manager to make advantageous rebalancing decisions. Advances in optimization software make it possible for the account manager to determine the weights of securities within a portfolio that optimize a desired objective, subject to investor-defined constraints.

After some terminology and a description of our standard methodology for analyzing strategies, this paper analyzes various DI strategies that are designed to help an account manager make effective decisions using the new flexibility of the DI framework, with an emphasis on TLH efficiency and dynamics around TE.

Definition of Terms

Common fundamental risk models provide a numerical measure for the *active variance* of a portfolio from a reference index. The active variance incorporates factor covariances and idiosyncratic risk to give a forward-looking estimate of return deviation. Like standard deviation, the TE of a portfolio against its benchmark is the square root of the active variance, which puts the TE back into return space. The condition of having a TE of 0 does not imply that the weights of the portfolio are exactly equal to the weights of securities that make up the reference index, but this condition does impose a tight equation connecting these two sets of weights. Permitting a nonzero TE gives much more potential for the account manager to achieve other benefits for the DI account.

Regarding TLH, there exists a few natural ways to measure the effectiveness of a strategy. The most direct metric is the overall taxes offset by realized capital loss or gains, for example due to investor withdrawals. Therefore, to quantify the effectiveness of a tax-aware portfolio, it can be helpful to construct a so-called shadow portfolio. One way to do so is to track a hypothetical portfolio which implements the index/strategy while conforming to certain accounting details such as rebalancing frequency, cash management, and transaction costs. With this tax-naïve portfolio as a reference, we define the *tax alpha* as the difference between the excess after-tax return between the two portfolios and the excess pre-tax return. More precisely, tax alpha is

$$\left(R_{p,at} - R_{s,at}\right) - \left(R_{p,pt} - R_{s,pt}\right),$$

where R_p is the return of the portfolio, R_s is the return of the shadow portfolio, and the subscripts *at* and *pt* refer to after-tax and pre-tax. This removes any pre-tax deviation in returns from the reference index, and therefore better captures the specific benefit of the TLH strategy. After-tax returns are calculated according to the USIPC After-Tax Performance Standards[1] which is used for GIPS compliance.

Standard Methodology for the Analysis

All data for this paper's analysis comes from decades' worth of stock price data in our Tax-Aware Strategy Rebalancer (TASR).[2]

The basic methodology behind the analysis of this white paper involves the following steps: (1) choose a DI reference index; (2) implement this strategy as a hypothetical portfolio; and (3) simulate its performance against historical data. Since TASR allows an analyst to adjust certain parameters involved in defining a DI strategy – from rebalancing triggers to parameters for portfolio optimization – we can simulate portfolio strategies through a variety of periods in recent history. Currently, there are 60 different control parameters. With TASR following terminology in the literature, each simulation is called a *backtest*. Each study involves varying a collection of control parameters and running one backtest on the hypothetical portfolio for each combination of parameters. We call the output of this process a *bundle* of backtests. To measure performance, TASR keeps track of both the portfolio and the shadow portfolio calculated by gross and net of fees; and by pre-liquidation, post-liquidation, and mark-to-liquidation.[3] In order to find trending behavior, we apply statistical methods to the resulting data of the backtest bundles to determine estimates of various properties. In the rest of this paper, we call this process our *standard methodology*.

Some analysts in the industry use a methodology that employs an artificial model of the market and studies the behavior of a DI strategy by running Monte Carlo simulations on how the DI portfolio would perform in that market. This methodology allows the analyst to control market parameters, such as expected return and volatility to describe how a strategy might fare under different market environment conditions. However, by their very construction, Monte Carlo simulations estimate a trending behavior that smooths out unique events. By using historical data, backtests in TASR demonstrate what potentially have happened using the strategy over a given time span, subject to more realistic assumptions. To use our standard methodology to explore how a strategy would perform under certain market conditions, we can run the strategy during a historical time frame that exemplifies a market environment against which we wish to test the strategy.

Though TASR gives the flexibility to change certain values, the analyses in this paper used a short-term capital gains rate of 40.8% and a long-term capital gains tax rate of 23.8%. For the backtests in this paper, we also used 1 basis point (bp) for transaction costs and subsequently analyzed the effect of other levels of transaction costs. See the relevant section in the appendix for more details.

[1] CFA Institute. "USIPC After-Tax Performance Standards." Accessed December 12, 2023. https://www.cfainstitute.org/en/about/governance/committees/usipc-after-tax-performance-standards. (This document coins the term "shadow portfolio" in the appendix, section H.

[2] See the document "Introduction to the Tax-Aware Strategy Rebalancer" for an in-depth description of TASR.

[3] These terms are defined in USIPC After-Tax Performance Standards, op. cit. 1.

Rebalancing Triggers

In practice, DI account managers are likely to employ a proprietary combination of frequency triggers and event triggers to consider rebalancing the portfolio. A frequency trigger simply means to check the portfolio on a daily, weekly, monthly, quarterly or annual basis as to whether there is some benefit to rebalance. Some common event triggers include:

* if the unrealized portfolio losses exceed some percent;

* if the tracking error increases beyond some fixed percent;

* if some security's price decreases by some fixed percent since it was purchased;

* if the cash in the portfolio increases beyond some fixed basis point level.

A portfolio manager could also choose not to define a regular rebalancing frequency but instead focus on a combination of event triggers along with forcing a rebalance after a fixed number of days since the last rebalance. All these combinations can lead to different outcomes for TLH effectiveness and TE. During time periods when the market is uniformly strong, few securities will offer TLH opportunities; even when this is not the case, the wash-sale rule and transaction costs may lead to diminished benefits from rebalancing too frequently. Because of the many possible combinations, this paper does not discuss the effectiveness of various combinations of rebalancing triggers but leaves the analysis for a future paper.

Tradeoff Coefficient

Role of the Tradeoff Coefficient

When rebalancing a DI portfolio in a tax-aware fashion, the account manager must balance two often opposing objectives: maximize the tax benefit and minimize the tracking error.

A first approach to optimization may simply involve imposing an upper bound on the tracking error and maximizing the tax benefit. Finance intuition makes it clear that if the maximum allowed TE is large enough, then the optimizer will be able to capture more possible tax losses. Empirical optimizations confirm this: there always exists a TE upper bound beyond which there is no additional benefit for TLH.

Figure 1. Risk-Reward Curve

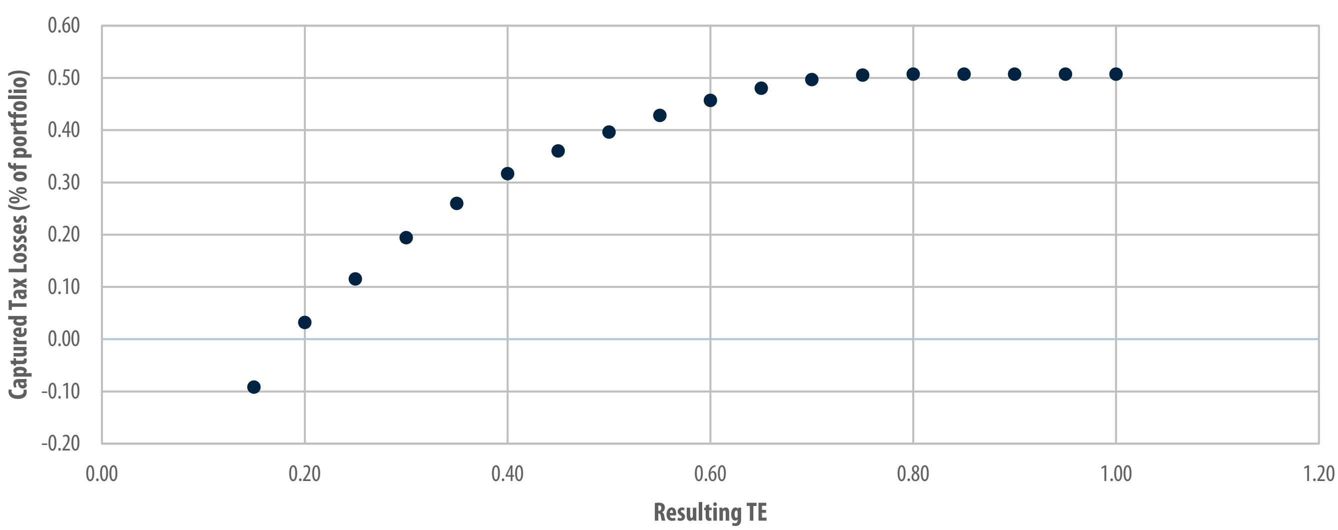

Figure 1 illustrates what happens for a particular rebalancing optimization. Each data point corresponds to a different imposed max TE bound, ranging from 15 bps to 1%, in increments of 5 bps. The (x, y) coordinates of each point are the ex-post TE and the captured tax losses resulting from the rebalance optimization. In this example, each resulting ex-post TE consistently goes up to the TE max bound, but the tax losses flattens out – once all harvesting opportunities have been captured. We call this curve a *risk-reward curve*. Every optimization displays a similar shape to this risk-reward curve: as tracking error increases, tax losses increase but is concave down and eventually flatten out completely. (Captured tax losses may be negative when a small tracking error forces the optimizer to realize capital gains.)

This general behavior indicates that for any individual rebalancing optimization, the benefits of allowing more TE cease at a certain point. This offers the opportunity to combine the desire to maximize tax credit and minimize the TE into a single objective function. To be more precise, the optimization objective has the form

$$objective = tax\ credit - transaction\ costs - \frac{active\ variance}{tradeoff\ coefficient},$$

where the tax credit is the realized tax loss times the investor's appropriate tax rate, the tradeoff coefficient is a user-defined constant, and the active variance is the square of TE. In other words, within the objective, we consider equally valuable a boost of one unit of tax losses (minus the transaction cost required to achieve it) and a unit decrease in active variance divided by the tradeoff coefficient. (Some authors refer to the inverse of the tradeoff coefficient as the *active variance aversion*.[4])

An important goal of this analysis is to equip the portfolio manager with tools to help decide what sort of behavior or choice of tradeoff coefficient might match an investor's preference. This section analyzes how the tradeoff coefficient determines the outcome of a particular rebalance in terms of tax credit and tracking error.

Analysis of the Tradeoff Coefficient

Analyzing the tradeoff coefficient is important and the reader should keep this in mind for the later analysis connecting TLH benefit with allowed TE. Furthermore, for an advisor using a similar approach to rebalance their portfolio in an optimal way, our analysis informs the advisor how to choose a value of the tradeoff coefficient to match their preference: whether to care more about tax credit and potentially allow a larger tracking error, or whether to aim for a smaller TE at the expense of potentially not capturing as much in tax losses. For simplicity, we will assume the transaction costs are negligible.

Another way to express the tradeoff coefficient is to consider equally valuable any pair $(tax\ credit, TE)$, where $tax\ credit - TE^2/(tradeoff\ coeff)$ is equal . This defines families of *indifference curves* of the form

$$H = \frac{TE^2}{tradeoff\ coefficient} + c,$$

where H stands for (harvested) tax credit, and the constant c corresponds to a value of the objective.

Figure (2) summarizes the relationship between the risk-reward curve and the indifference curves. Each indifference curve is a parabola with vertex c and where the tradeoff coefficient controls the parabola's convexity – the larger the tradeoff coefficient, the wider the parabola opens. Fixing the tradeoff coefficient but varying c defines a family of indifference curves. Figure (2a) and Figure (2b) respectively have a tradeoff coefficient of 0.03 and 0.1, showing three indifference curves in each family.

[4] For example, see the appendix of Joseph Liberman, Stanley Krasner, Nathan Sosner, and Pedro Freitas, *Beyond Direct Indexing: Dynamic Direct Long-Short Investing*, The Journal of Beta Investment Strategies, vol. 14, Issue 3, Fall 2023.

Figure 2. Risk-reward curves with indifference curves

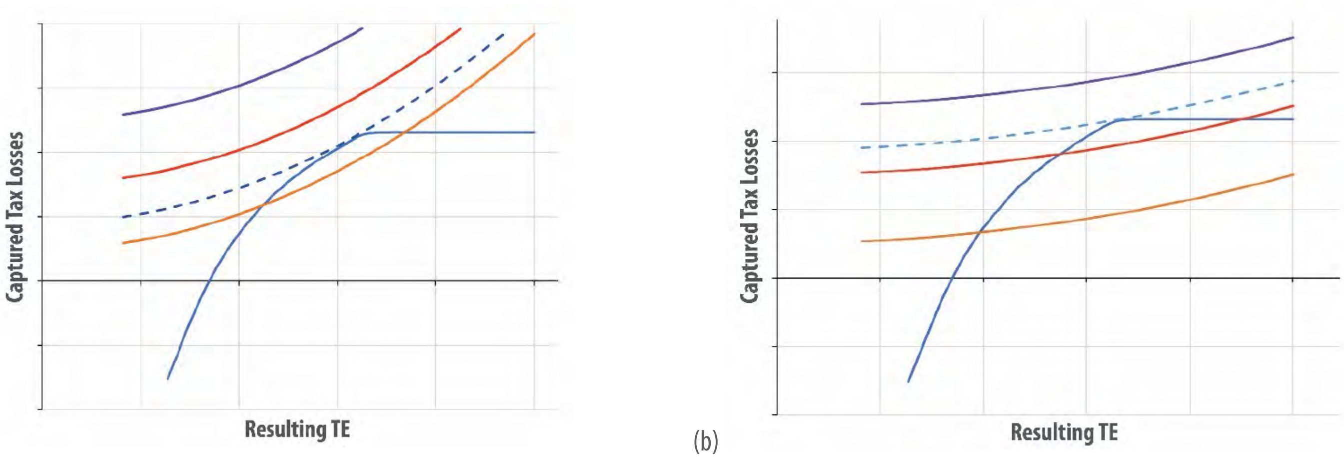

(a) (b)

For a given value of the tradeoff coefficient, the result corresponds to an indifference curve that is tangent to the risk-reward curve. The optimized value of the objective function is the c value for the specific indifference curve in the given family. [5]

The primary challenge in using a tradeoff coefficient in the objective function lies in understanding what range of values might be useful. Furthermore, the concept of "useful" depends on the investor's preference. Consequently, we need to understand how a specific value of the tradeoff coefficient affects the balance of maximizing tax credit while minimizing tracking error.

Along a given indifference curve, the constant c is fixed. Taking rates of Equation (1) for any indifference curve, we get

$$\Delta H = \frac{2 \times TE \times \Delta TE}{tradeoff\ coefficient}.$$

As stated earlier, when the objective function is maximized, the corresponding indifference curve must be tangent to the risk-reward curve. The optimum point will occur at or to the left of the point on the risk-reward curve where it flattens out. At this optimum point, the slopes of both curves are equal to

$$\frac{\Delta H}{\Delta TE} = \frac{2 \times TE}{tradeoff\ coefficient}.$$

Furthermore, the indifference curves are convex, while the risk-reward curve is concave down. Hence, as the TE increases, the slopes of the indifference curves increase, while the slope of the risk-reward curve decreases. To capture H in TLH, we need a tradeoff coefficient greater than or equal to

$$\frac{2 \times TE}{(\Delta H / \Delta TE)},$$

where (H, TE) is a point on the risk-reward curve and where ${}^{\Delta H}/_{\Delta TE}$ is a rate of change (left-slope) at (H, TE) on the risk-reward curve.

[5] If we do not ignore transaction costs, the constant c in Equation (1) corresponds to a fixed level of the objective function plus the transaction costs. Allowing for the constant c to fall in some very small range corresponds to a very narrow family of indifference curve parabolas.

This observation inspires a strategy to study the effectiveness of the tradeoff coefficient. By looking at many different backtests, each of which involves only one rebalancing event, but varying the value of the maximum TE bound, we get the risk-reward curves for many different optimization situations. For each risk-reward curve, we find the maximum amount of TLH possible and then use the formula in Equation (2) to determine the minimum required value of the tradeoff coefficient to capture any percentage (e.g., 99%, 95%, 90%) of this maximum harvestable tax loss. Applying this method to many different risk-reward curves allows us to calculate an expected value and standard deviation of the tradeoff coefficient required to capture some p percent of all harvestable losses.

To implement this study using TASR, we ran a bundle of backtests with the following parameters:

Benchmark: A US Large Cap Index
Rebalance frequency: Monthly
Maximum TE: 0.002 / 0.0025 / . . . / 0.03
Fees: 30 bps
Tradeoff Coefficient: NULL
Time span: 1 month; starting every quarter from 2002 to 2023.

This resulted in 88 different months (optimizations) and 57 different values for the max TE bound. By not giving a value to the tradeoff coefficient, TASR only optimizes the tax credit.

Figure (3) illustrates a typical risk-reward curve resulting from one backtest (one optimization).

Figure 3. Risk-reward curve for April 2006.

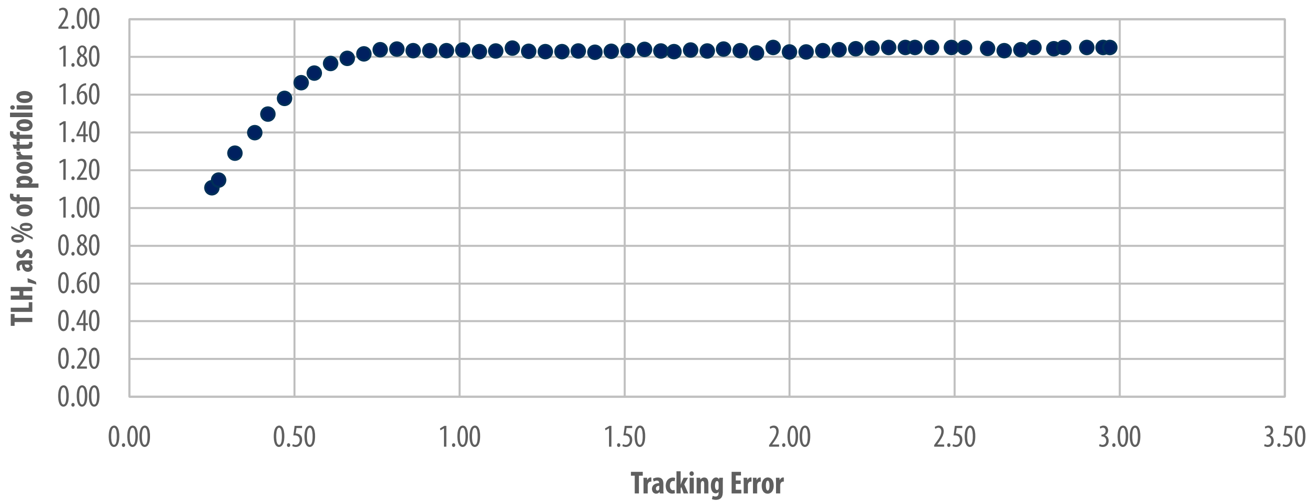

For this example, the maximum TLH amount is $H = 0.01776$, and this occurs first at $TE = 0.0115$. From this data, we can estimate the required tradeoff coefficient to reach 99%, 95%, and 90% of the maximum TLH amount as

Percent of max TLH	99%	95%	90%
Estimated tradeoff coefficient	0.07859	0.01307	0.01305

(In this example, harvested tax losses nearly reach 2% of the portfolio value. However, we emphasize that no DI account will see this amount of harvesting every month. This rebalance optimization corresponds to what would have occurred in the first month of a portfolio that starts from cash. Subsequent months will naturally offer fewer opportunities for TLH than this first month. Nevertheless, this approach is valuable to estimate the role of the tradeoff coefficient at the single-optimization level.)

Figure (4) illustrates the maximum, the minimum, and the average of risk-reward curves for various months. Note that the minimum curve, which corresponds to July 2020, offered so few opportunities that some data points are spurious.

Figure 4. Extreme and average portfolio risk-reward curves.

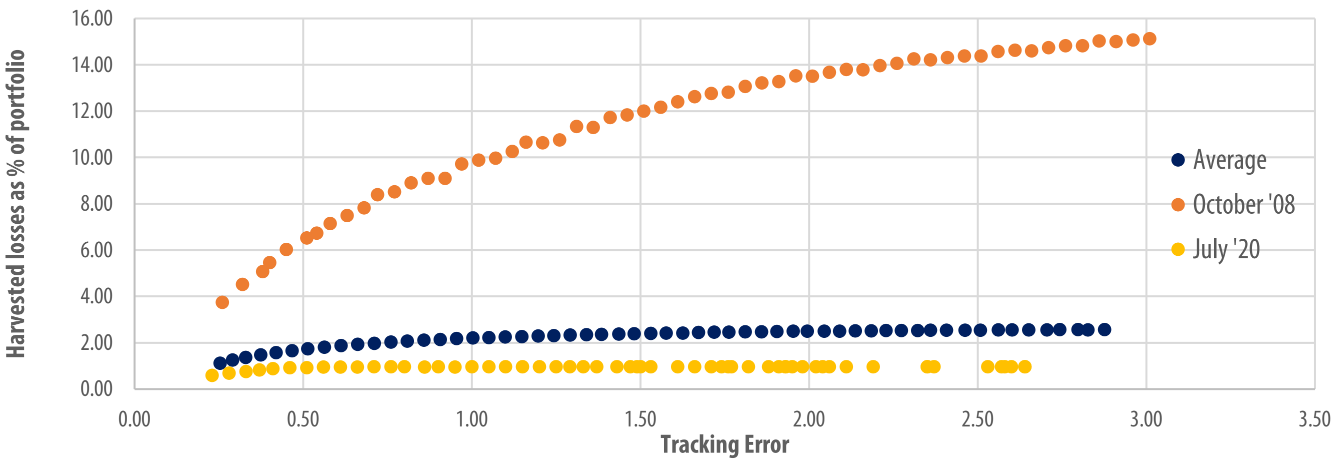

Because of vastly different market conditions, the risk-reward curves differed considerably in magnitudes, though not in overall shape. Some months flattened out as early as 0.5% TE while others required out to 4.5% (October 2008). Furthermore, two of the 88 months exhibited so few harvesting opportunities as to render the risk-reward curves unusable to estimate the required tradeoff coefficient. After discarding these two exceptions, we used the 86 remaining sample months to calculate the required tradeoff coefficient to harvest p percent of the total harvestable amount, q percent of the time. Figure (5) below depicts these estimates.

Figure 5. Capture TLH percent at tradeoff coefficient.

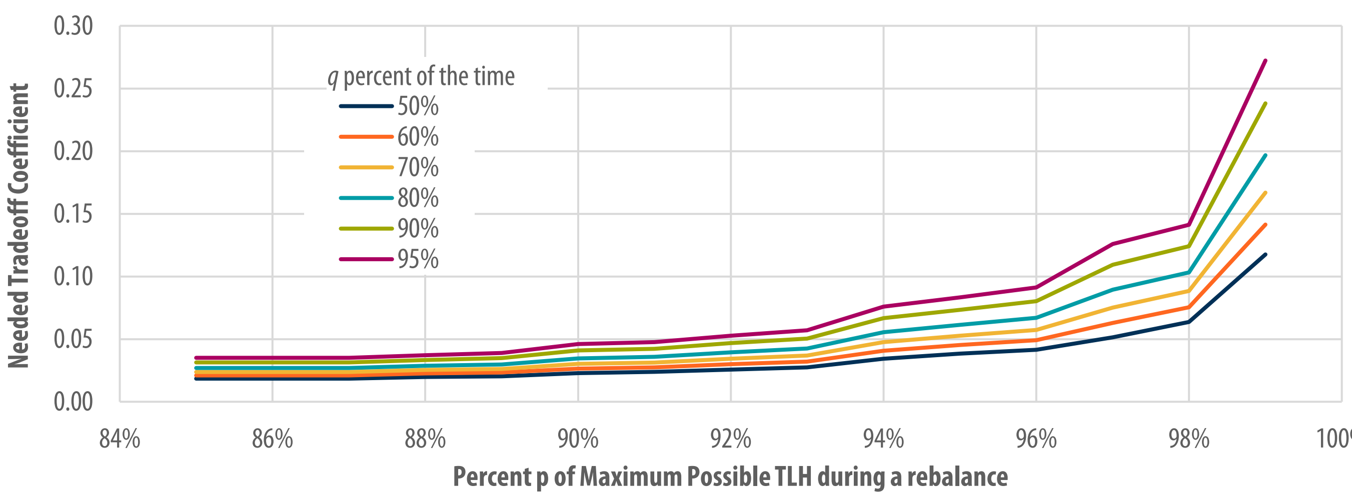

We notice that to capture 98% of the harvestable tax losses 95% of the time, we need a tradeoff coefficient of at least 0.124. Furthermore, at that same tradeoff coefficient, we could capture 99% of harvestable tax losses about 55% of the time. We can also say that with a tradeoff coefficient of 0.05, we already capture 90% of the harvestable tax losses nearly all the time. The data appears to support the hypothesis that a tradeoff coefficient above 0.2 seems generally unnecessary in that it can offer only tiny fractions of additional TLH, while the TE can continue to rise.

The analysis in this subsection appears to transcend the specific portfolio. Repeating this work with a different strategy (involving a larger name count) led to nearly identical results. Ostensibly, an idiosyncratic DI portfolio might deviate from this trend, but for brevity, this paper does not explore this.

Analysis of the Tradeoff Coefficient Over the Long-Term

The previous analysis of the tradeoff coefficient focused on the level of a single optimization. A priori, the relationship between tradeoff coefficient, TLH, and TE might change over the life of a DI account.

In this study, we chose to let the DI account proceed for one year, rebalancing the account monthly, in a traditional sense. Furthermore, not to let the results of the previous subsection bias the scope of our investigation, we analyzed the account behavior while varying the tradeoff coefficient through a wider domain. More precisely, we used the following parameters.

Benchmark: A US Large Cap Index
Rebalance frequency: Monthly
Cash in portfolio: less than 75 bps (otherwise trigger a rebalance)
Maximum tracking error: 1.5%
Fees: 30 bps
Tradeoff coefficient: many values between 0 and 0.5
Time span: 1 year; for each year between 2005 to 2022.

As an example, Figure (6) shows one scatter plot (a) relating the overall tax alpha that would have occurred in calendar year 2010 with a tradeoff coefficient between 0 and 0.5, and a second plot (b) showing the similar data zooming in on the tradeoff coefficient between 0 and 0.06.

Figure 6. 2010 Tax alpha versus tradeoff coefficient.

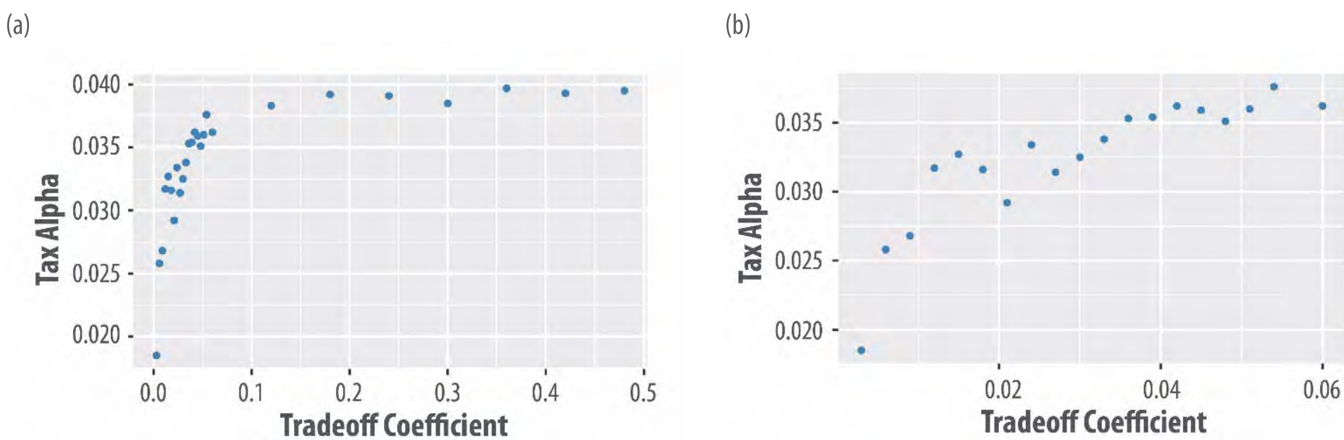

All the scatter plots of tax alpha versus tradeoff coefficient for every year between 2005 and 2022 exhibit a similar behavior: the data rise very quickly near 0 tradeoff, but then flatten out. These observations dovetail well with the intuition developed by our analysis at the single optimization level in the previous subsection. Furthermore, it is not surprising that the jaggedness in Figure (6b) might occur with a DI account observed over multiple rebalances.

Figure (7) displays the first quartile, the third quartile, and the mean of the data points through years 2005 through 2022, thereby capturing a trend of central behavior in this relationship between tradeoff coefficient and TLH.

Figure 7. Top and lower quartiles, and average of

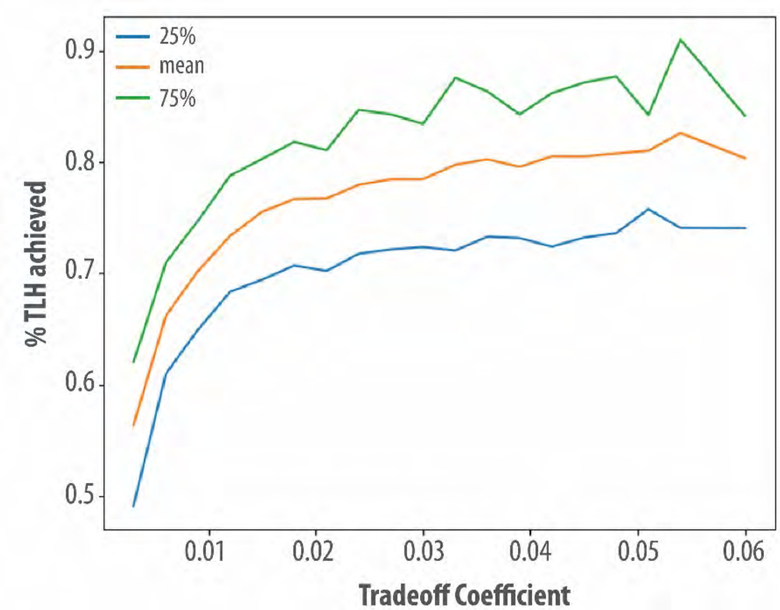

The reader should note that, barring some jaggedness due to multiple rebalances, Figure (7) is extremely similar to Figure (5), except reflected across the inverse. Consequently, though our analyses in these two subsections approach the role of the tradeoff coefficient from opposite directions – the level of a single optimization and the level of a backtested DI account over time – the conclusion about which range of values leads to what types of behavior remains generally the same. In other words, passing from a single optimization to a sequence of rebalance optimizations does not systematically change how the tradeoff coefficient affects the percent of achievable TLH.

Tradeoff Coefficient Conclusion

Obviously, it is possible for an investor to manage a DI account by fixing some TE bound and simply maximizing the tax credit. However, doing so fails to capture how TLH and TE play off each other and could cause the investor to potentially miss some significant opportunities. If an analyst or advisor wishes to take advantage of the dynamic nature of the objective function defined at the beginning of this section, we believe he or she will need to choose a strategy that involves setting a value for the tradeoff coefficient.

Long-Term Effectiveness for TLH

Aggregated Analysis

Since TLH stands as a common reason for investing according to a DI framework, we believe it is critical for both the investor and financial advisor to develop some intuition about TLH over time. More precisely, we would like to know what the active tax benefit is as a function of time. For this study, we will define the tax benefit as the cumulative tax losses harvested over time. We use TASR with the following parameters.

Benchmark: A US Large Cap Index
Initial investment: $1,000,000
Rebalance frequency: Weekly
Max cash in portfolio: less than 75 bps (otherwise trigger a rebalance)
Fees: 30 bps
Time span: 10 years; starting on the last trading day of each month between 4/30/2001 and 11/30/2013.
Min Transaction Size: 99.5 / 99.75 / 100 / 100.25 / 100.5

For the 760 backtests produced by these parameters, we record the cumulative tax losses harvested after 1 year, after 2 years, and so on up to 10 years. Figure (8) depicts the resulting data of cumulative losses measured as a percentage of initial investment with a boxplot as related to the number of years since portfolio inception. In this boxplot, the orange box above each independent variable forms the interquartile range (IQR: the outputs between the first and third quartile), the horizontal blue line is the median, the upper whisker reaches to the largest value within 1.5xIQR above the 75% quartile, similarly for the lower whisker, and the other dots are outliers beyond the whiskers. The red diamonds indicate the mean of the data for any given year.

Figure 8. Aggregate data of cumulative losses

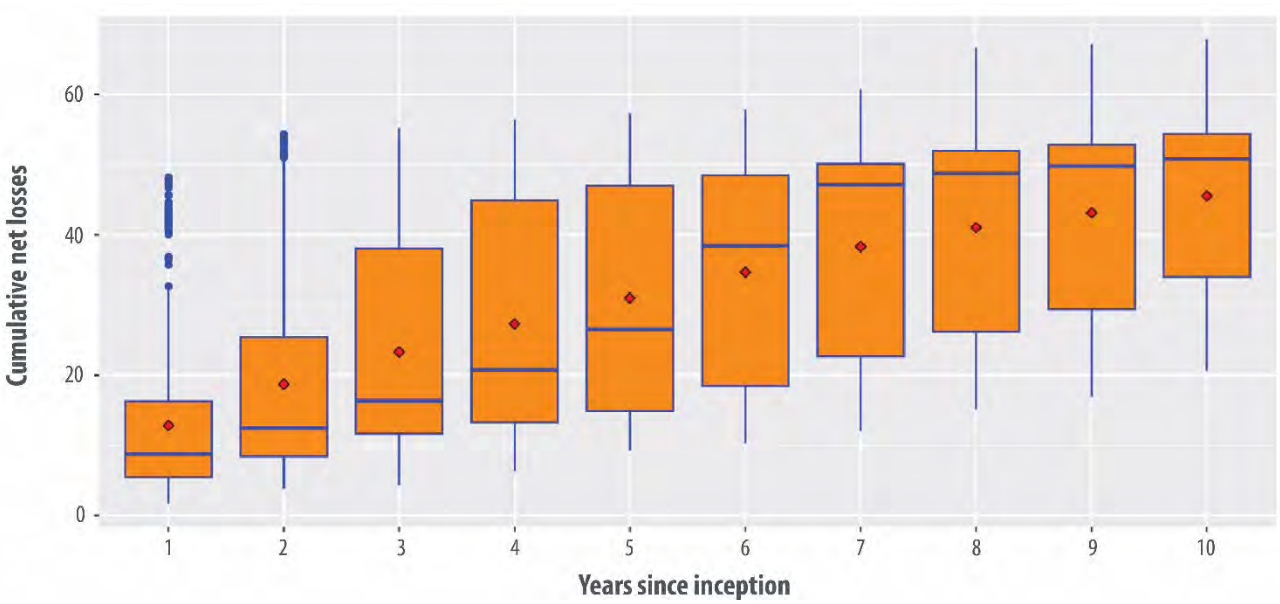

The large interquartile ranges (IQR) indicate that the data exhibits a considerable spread in the aggregated behavior of the TLH function over time. The average appears to follow a curve that flattens toward an asymptote. This phenomenon is called *ossification*. We say that the portfolio "ossifies over time" to mean that after a while there is a near 0 probability of ever harvesting further losses. The smaller third quartile boxes corresponding to years 7 through 10 confirm how tightly this ossification phenomenon occurs.

We should expect ossification for the simple reason that an investor can only claim losses up to the cost basis of an investment. From a theoretical perspective, the maximum possible TLH is 100%. Though we will soon show why this reason is problematic, many analysts might focus on the Figure (8) data of the yearly TLH average, notice that the data points fit well to a flattening exponential function, and use that to quantify ossification levels. Suppose we call y the cumulative losses harvested and t the time since account inception and suppose we use the model $y(t) = A(1 - e^{-kt})$ with the assumption that k is positive,[6] then the averages by year that occur in Figure (8) fits to

Figure 9. Exponential model fit

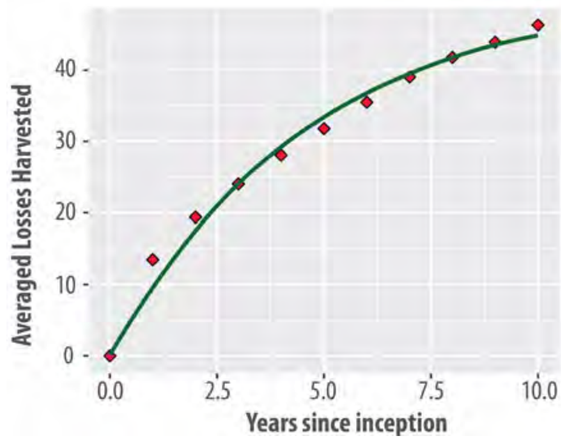

$$y(t) = 51.0135(1 - e^{-0.2127\,t}).$$

This nonlinear regression has a standard error of 1.646, which indicates a very good fit, given that $\bar{y} = 29.4$. The estimated asymptote of averaged data by year in Figure (9) is around $y = A = 51$, which means around 51% of the invested amount. Setting a level at $r = 95\%$ or some other value, we could further define the r-*ossification level* as the point in time when the captured tax losses reach $r = 95\%$ of the asymptote amount. In this model, this occurs when

$$t = -\frac{1}{k}\ln(1 - r) = 14 \text{ years.}$$

As mathematically elegant as the reasoning in the above paragraph might seem, it is unhelpfully simplistic. The very large IQR, even for later years in Figure (8), underscores that there can be a very wide spread of behavior. Indeed, because of the wide standard deviation, it would be irresponsible to tell a DI account holder that they could expect this average. As we will see in the next section, the ossification phenomenon is not a simple event and TLH effectiveness depends critically on market conditions.

[6] This is one of the simplest, in the sense of having the fewest parameters, and most theoretically natural mathematical models that goes through (0,0) and has a horizontal asymptote at $y = A$.

TLH Effectiveness in Detail

By partitioning our backtests by start year, we discover that the uniformity suggested by the aggregated data is not a realistic model for a specific DI account. TLH depends on tax loss harvesting opportunities and the cost basis of the current portfolio. Harvesting opportunities depend significantly on market conditions, which in turn depend on time. Consequently, we believe it is valuable to simulate the same relationship of cumulative realized losses, but for accounts starting in the same year, which means experiencing similar market conditions. Figure (10a-c) de-aggregates the data from Figure (8) by restricting to specific starting years. Each diagram to the right summarizes 60 (10 year) backtests.

As a first observation, we note that the IQR at each input data value has shrunk considerably. This supports the intuition that during periods of similar market conditions, the cumulative tax loss functions behave similarly.

The market downturn of 2008 figures prominently in the TLH function corresponding to backtests starting in 2003 and 2007. The ossification phenomenon remains particularly obvious for these two start years, because the large amount of TLH in 2008 moves the TLH yield closer to an ossification point of around 60%. DI strategy started opening in 2009 did not experience a recession during the first 10 years. For such accounts, we do not observe the cumulative losses nearing any ossification level within that 10-year time span and they do not even reach 30%.

The three starting years of 2003, 2007, and 2009 represent three distinct scenarios: a bear market occurring “later,” a bear market occurring “earlier” (i.e., shortly after account inception), and bull market conditions throughout. These three situations warrant their own analysis, especially as compared to holding an ETF of the same stocks and weights and harvesting tax losses.

It might interest the reader to know that during the 10-year backtests starting in 2003, 2007, and 2009, the annualized return of the benchmark was 7.42%, 7.31%, and 13.97% respectively. These percentages illustrate the intuitive property that the cumulative tax loss harvests are inversely correlated with benchmark returns.

Figure 10. Cumulative losses, disaggregated by start year

(a) 2003

(b) 2007

(c) 2009

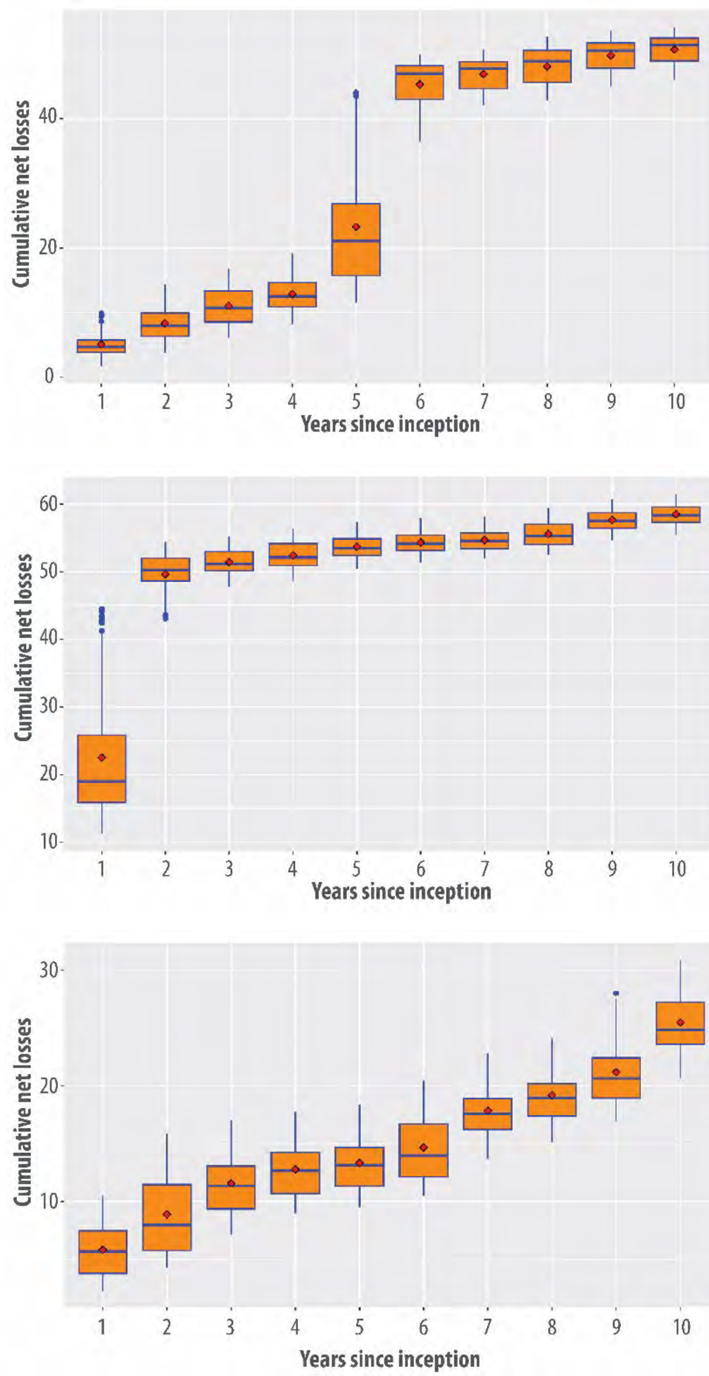

To further understand the effectiveness of direct indexing over the longer term, we compared the performance obtained above to that of an account holding a single ETF in absence of the wash sale rule. More specifically, we chose a popular US large-cap ETF with close overlap to the index used above and simulated one backtest per week from 2001 to 2013, with each backtest running for 10 years. We checked weekly to reinvest dividends and harvest losses (by selling and immediately re-buying). Figure (11) shows cumulative net losses for the ETF and DI approaches, with values smoothed by taking averages over a given year.

Figure 11. Comparison of TLH between a DI account and an equivalent ETF

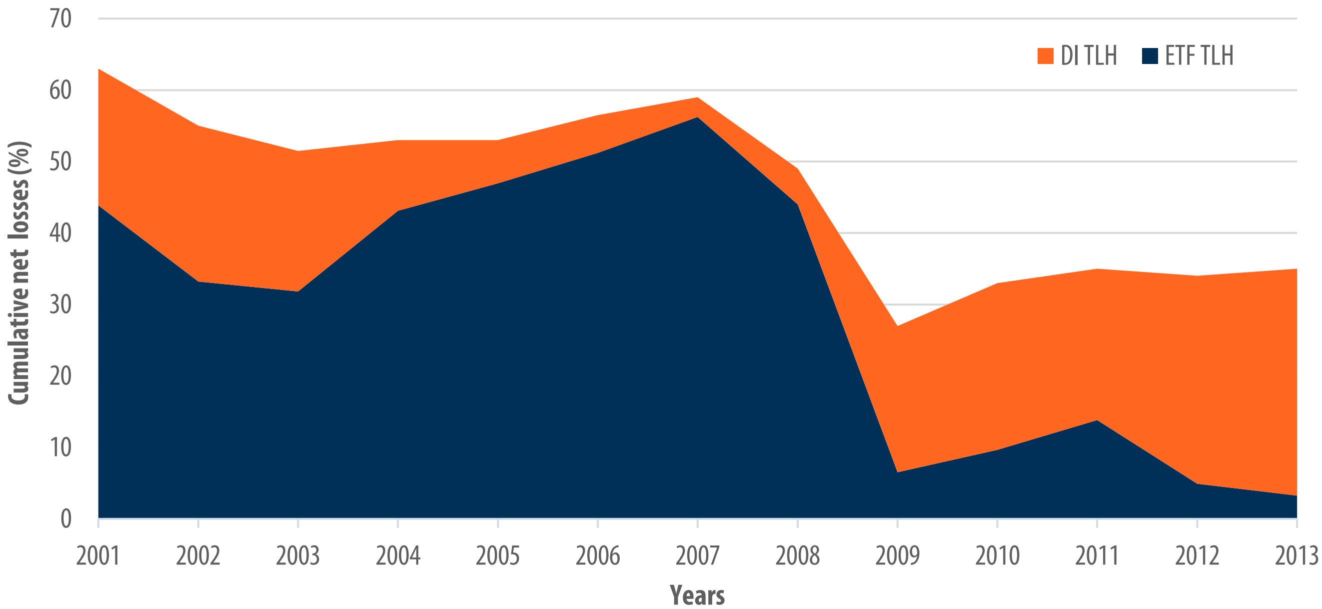

Figure (11) shows the DI account performs better than an ETF regarding TLH, and particularly when the market is strong throughout or when a bear market occurs later in the life of the account. We summarize this important result below.

- **Bear market occurring later**: Figure (10a) illustrates what happens for a tax-aware DI strategy: the cumulative realized tax losses approach 50-55% of the initial portfolio investment. However, for an equivalent ETF, if the market starts strong (e.g., for the first 5 years) it is likely that the overall value of the ETF will not dip below the initial investment during the first 5 years and therefore not lead to any realized tax losses. If a bear market occurs around the 5-year mark, the ETF will not offer any realizable tax losses until the ETF value dips below the initial investment. So, assuming the value of the ETF grew during those first 5 years, it is not likely this will occur. Consequently, in this situation, it is not likely that an equivalent ETF would offer any realized tax losses.
- **Bear market occurring earlier**: Figure (10b) shows a tax-aware DI account quickly reaching realized losses of 50%, and then subsequently moving slowly toward an ossification level of 60%. For an equivalent ETF, since the bear market occurs shortly after initial investment, the value of the ETF will not have grown much before the downturn and hence it is likely that the ETF will drop below the initial investment. Hence, the investor would be able to claim a loss on this investment, perhaps even a large percentage of what a tax-aware DI account could claim during the first year.
- **Bull market conditions throughout**: Figure (10c) displays how a tax-aware DI account may not be close to ossification; for these 10-year backtests the TLH yield does not even reach 30% of initial investment. Yet, if market conditions remain strong as in Figure (11) from 2009 to 2013, it is likely that the equivalent ETF will simply grow without ever dropping below the original investment. Consequently, we do not expect many realizable tax losses for the ETF, while a DI account may continue harvesting tax losses.

In all three situations, the DI account fares better than an equivalent ETF regarding TLH effectiveness. However, if a bear market occurs "later" (i.e., a few years after account inception) the tax-aware DI account will fare the best over time as compared to an equivalent ETF, realizing 50% losses over 10 years. Though we do not approach ossification over 10 years, the DI account holder will harvest plenty of tax losses whereas someone holding an equivalent ETF would hardly harvest any. Finally, when a bear market occurs within the first year of a DI account's inception: though a tax-aware DI account will still harvest more losses than an ETF, that ETF may decline in value and hence allow the holder to harvest many of those losses.

Figure 12. Cumulative net losses for DI account starting in 2013.

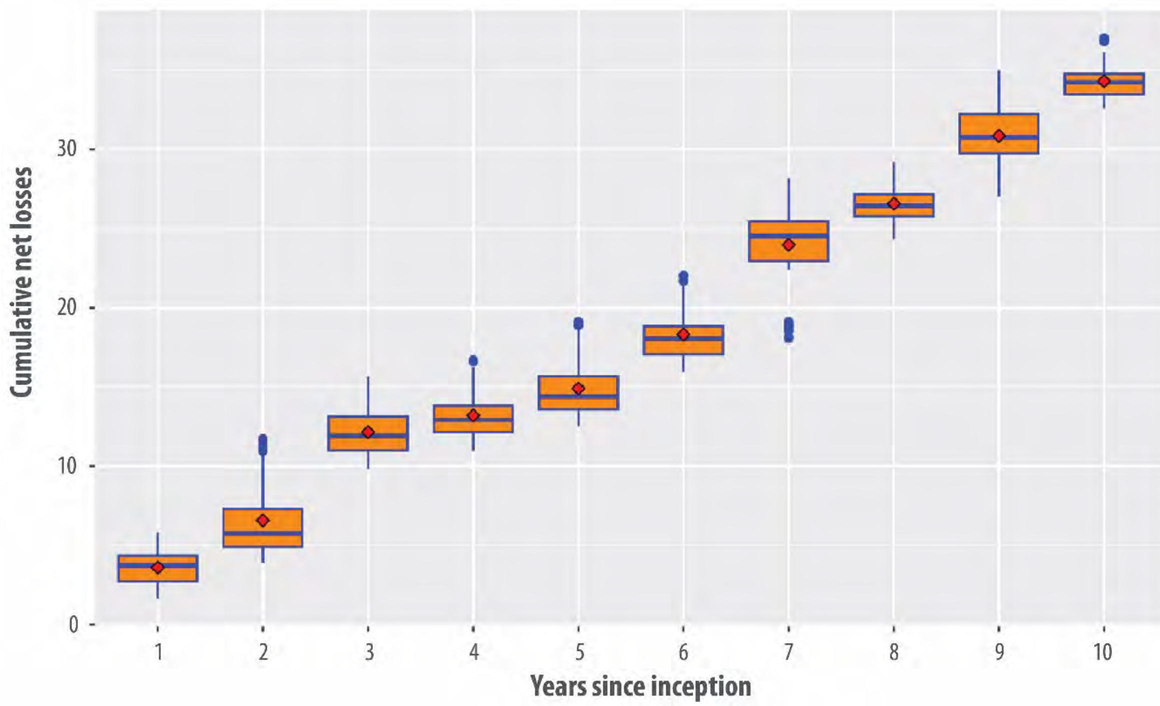

Figure (10a) above offers a visual of what may occur when there is a sharp downturn (the crash of 2008) occurring 5 years after account inception. Figure (12) illustrates simulation against historical data for a DI strategy started in 2013. Beginning around year 7, this involves another bear market that was not as sharp, but instead more prolonged (2020-2022).

Figure (12) shows a portfolio that appears to ossify early between years 3 to 5 but then takes advantage of more TLH opportunities after 2019.

A key takeaway of this analysis of backtest bundles disaggregated by start year is that though Figure (8) depicts the average behavior of the realized tax losses over time, any single investor is likely to encounter a more individual experience, affected dramatically by changing market conditions. The ossification phenomenon remains, but how quickly the DI account reaches ossification and what that level might be depends dramatically on market conditions. Furthermore, our analysis identifies at least three different scenarios where a DI account exhibits different TLH behavior as compared to an equivalent ETF. Comparing Figure (10a) with Figure (12), we also notice a finer difference in the "Bear market occurring later" scenario between a sharp and short downturn and a less sharp but more protracted downturn in 2020-2022.

TLH Function General Model

To go beyond a data-driven narrative of TLH behavior toward a more functional understanding, we use a general model that could forecast a DI account's TLH function over time. Though not yet a complete model, it appeases to our empirical observations and offers explanations for phenomena we see in this section, but also for results in the next sections.

Our observations support the following general model. Denote by y the amount of tax loss harvested (measured in dollars) and t the time since portfolio inception (measured in years). Call M the *loss carrying capacity* of a tax-aware DI account. If the account involves a one-time investment P, then the loss carrying capacity is proportional to P, so that $M = cP$, where c depends exclusively on market conditions, not on y. The principle behind the model is that the rate of tax loss harvesting is proportional to the remaining loss capacity. More specifically, with derivatives, we have

$$\frac{dy}{dt} = k(cP - y),$$

where k does not depend on y or P. We call c the *relative loss carrying capacity*. The coefficient k reflects the concept of TLH opportunity. This equation is a general model because the coefficients k and c are positive values that depend on market conditions and portfolio management decisions but are independent of the tax losses harvested so far. At the outset of a DI account, no losses have been harvested and so the model involves the initial condition of $y(0) = 0$.

As dependent on the market, the coefficients k and c might be nonconstant functions of time. Because the rate of harvesting will be positive as long as y is below the carrying capacity, we have $k > 0$. Clearly, we also have $0 < c \leq 1$, but we observe empirically that the relative loss carrying capacity c is bounded above by some $c_{max} < 1$. None of our backtests ever exceed a c value of 0.7.

With this model, y increases more the further away it is from the loss carrying capacity cP. Consequently, since at the inception of the account when $y(0) = 0$, if cP were constant, then the first year would produce the most effect on TLH. This explains in a natural way the conclusion by authors that a DI account tends to harvest most of its losses in the first year. However, both our empirical observations and this model give considerable nuance to this thesis. If k and c remain constant, then the result of this differential equation is an increasing and flattening exponential curve as in Figures 5 and 6. Since it is more likely that k and c depend on time, we obtain far more nuanced time development.

In all cases, since the cumulative TLH function $y(t)$ starts at 0 and hence below cP, according to this model $y(t)$ always increases toward cP, without ever surpassing it. Explicitly expressing c as a function of time, this model implies that the asymptotic (long-term) behavior of $y(t)$ is below $c(t)P$, and hence below the constant $c_{max}P$.

As an example of using this model, we wanted to naturally simulate a DI account opened in 2003. We imagined a TLH opportunity factor k of 0.6 but jumps to 1.0 only for $5 \leq t \leq 6$; whereas $c = 20$ for $t \leq 5$ but $= 55$ for $t > 5$. A differential equation solver produces the curve shown in Figure (13), which fits the average backtest data in Figure (10a) well.

These observations suggest that the TLH function of a DI account behaves according to a stochastic differential equation with a drift coefficient of the form $k(cP - y)$, with k and c as explained above. Deeper analysis on how k and c depend on market conditions (and therefore on time) goes beyond the scope of this white paper.

Figure 13. Differential equation model

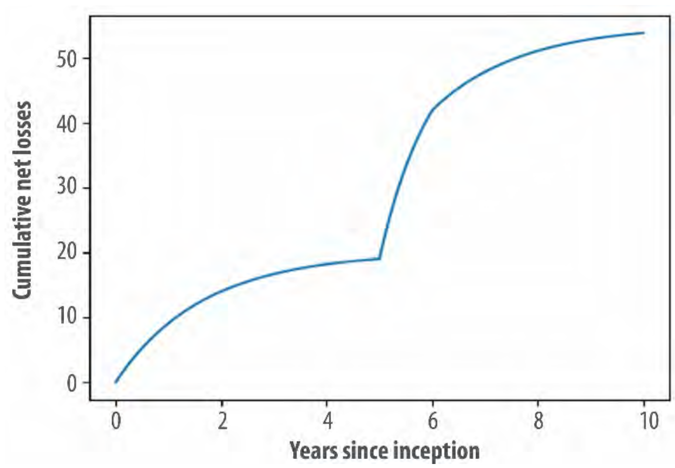

Mitigating Ossification with Cash Infusions

Aggregated Analysis

The previous section presented the phenomenon of ossification of DI portfolios, though we also qualified it for different market scenarios. However, the prospect of a DI account losing its TLH effectiveness leads us to consider some strategies to mitigate the ossification phenomenon.

One such strategy considers complementing the static investment with cash infusions. To study the effect of cash infusions, we used the same bundle of backtests in TASR as in the previous section but now simulating the optimized portfolio rebalancing over 10 years while also adding regular cash infusions at 5 different levels. We used these parameters:

Benchmark: A US Large Cap Index
Initial investment: $1,000,000
Rebalance frequency: Weekly
Max cash in portfolio: less than 75 bps (otherwise trigger a rebalance)
Fees: 30 bps
Time span: 10 years; starting on the last trading day of each month between 4/30/2001 and 11/30/2013.
Min Transaction Size: fixed at $100
Cash Infusion: $5,000 / $10,000 / $25,000 / $50,000 / $100,000

This produced 760 backtests:152 starting dates times 5 levels of cash infusion. Figure (14) depicts a waterfall chart of the average (across all backtests) realized tax losses, as a function of cash infusions.

In the 5 levels of cash infusions, the average realized tax losses per $1k of cash infusion remains fairly even. The following chart provides the calculation.

Figure 14. Average realized tax losses by annual cash infusions

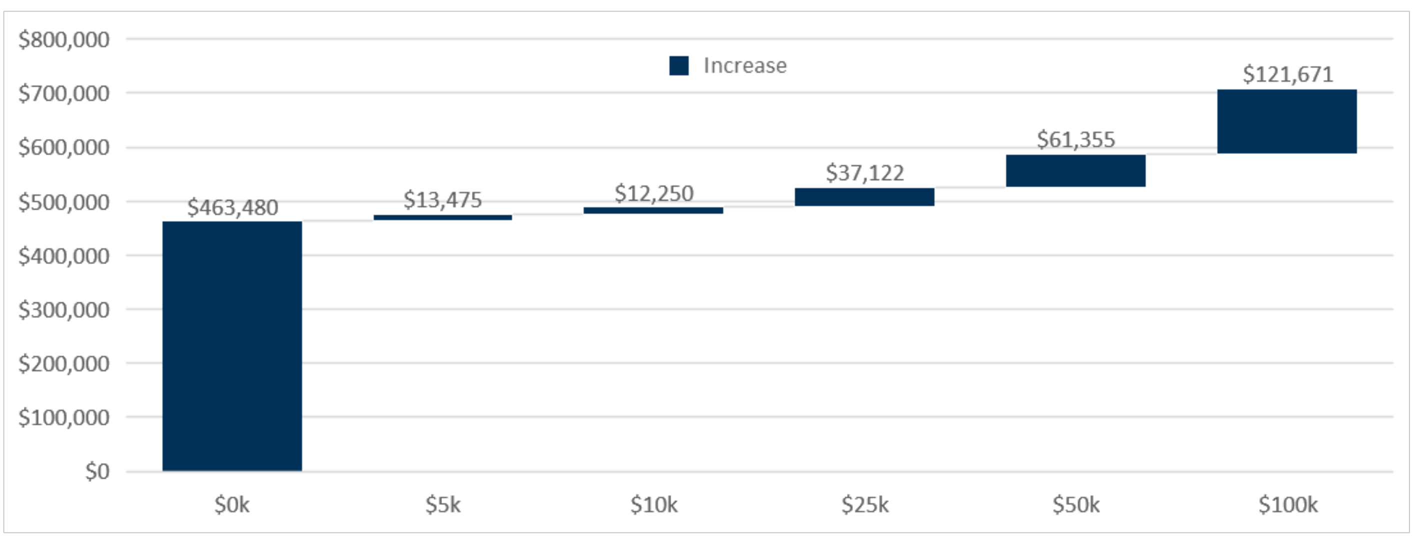

Regular infusion amount	$5,000	$10,000	$25,000	$50,000	$100,000
Realized losses / $1k cash infusion	$2,695	$2,572	$2,513	$2,484	$2,458

This data, averaged over all backtests, shows that for about every $1,000 of yearly cash infusion, potentially around $2,500 of additional realized tax losses over the course of 10 years. However, this analysis involves data occurring at a single point in time; we provide a forward-looking analysis below to estimate the benefit of cash infusions in the longer term.

Regular cash infusions do not only benefit the realized tax losses, but also the ex-ante TE. In the bundle of backtests that generated this data, increasing regular cash infusion amounts decreased the ending ex-ante TE for the portfolio. Figure (15) shows specific amounts for this effect.

Figure 15. Average ending ex-ante TE in terms of regular cash infusion

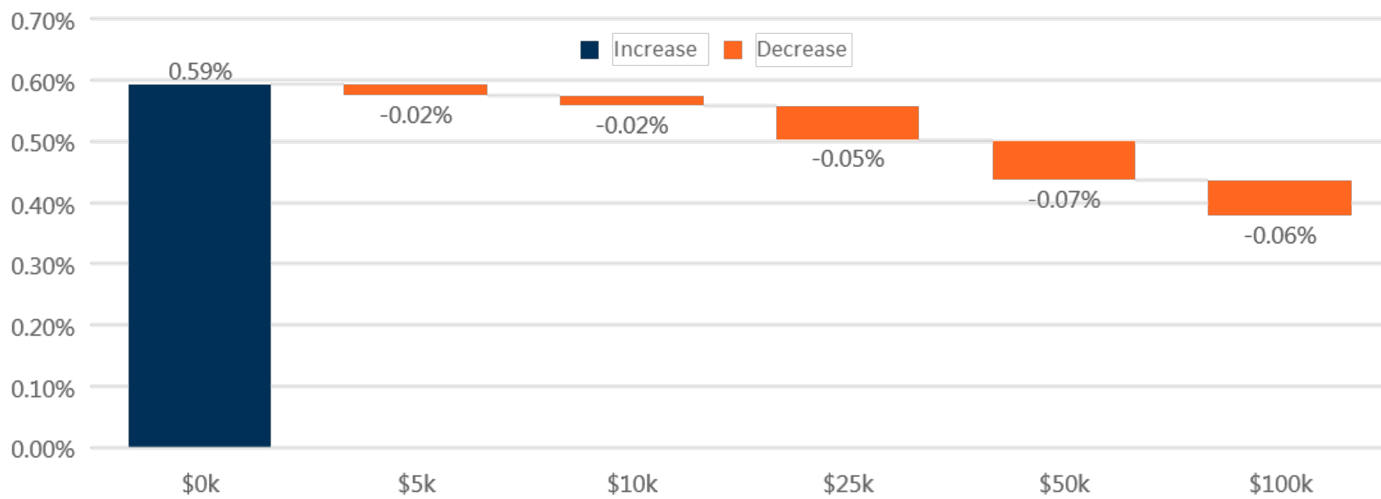

Regular cash infusions not only benefit the TLH but also benefit the portfolio's TE, giving a DI account with cash infusions a double benefit. It is likely that the benefit to TLH could be increased at the expense of ending-TE by choosing a different tradeoff coefficient, but we do not explore this further here.

Analyzing cumulative net losses over time for all levels of cash infusion allows a head-to-head comparison of trends obtained in the previous section on ossification with trends of what would occur with cash infusions. As in the previous section, we begin by aggregating all the data, without distinguishing between start years, which in turn means, without distinguishing between market conditions. In parallel to Figure (10), Figure (16) depicts this aggregated data, but split into different levels of cash infusion. The vertical axis gives the cumulative tax losses captured as a percent of the initial cost basis.

Figure 16. Cumulative TLH (as percent of initial) at different levels of cash infusions

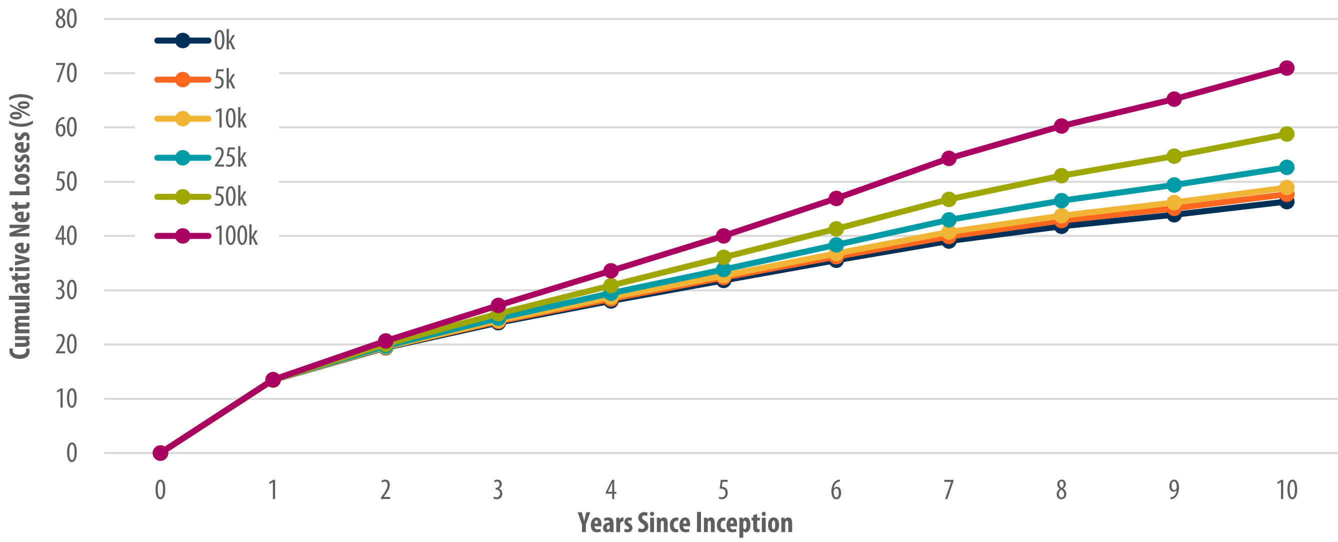

Instead of a trending curve that tends toward a horizontal line, the function of TLH as a function of time appears to tend toward a slant line, i.e., a line with nonzero slope. Estimating this trending asymptotic slope on the last 4 years gives these results.

I: Infusion Level ($1000)	5	10	25	50	100
m: Asymptotic Slope (% of initial /yr)	2.557	2.703	3.194	3.967	5.488
$\Delta m/\Delta I$		0.0291	0.0327	0.0309	0.0305

The third line of the table calculates the rate of change of the asymptotic slope with respect to levels of cash infusion ($\Delta m/\Delta I$). Since this is so consistent (averaging 0.0308 ± 0.0013), we conclude that the asymptotic slant line that describes the long-term behavior of a DI account with cash infusions has a nonzero slope that is a fixed multiple of the cash infusion level. Using forward-looking linear regression estimates, these figures replace our previous observations occurring at the 10-year snapshot. In particular, the information shows that in the long run, a DI account with yearly cash infusions offers around $310 of realized tax losses per $1,000 yearly cash infusion per year since account inception. This ratio quantifies to what extent yearly cash infusions mitigate the ossification phenomenon.

Effectiveness of Cash Infusions in Detail

The previous section underscored that disaggregating simulations against historical data into different start years displays significant variety for realizing tax losses due to year-dependent market conditions. We should expect a variety of experience with a DI-account that also has regular cash infusions.

Figure (17) below imitates Figure (10) but shows the different levels of cash infusions for the start years of 2003, 2007, and 2009.

Figure 17. Average cumulative net losses with regular cash infusion

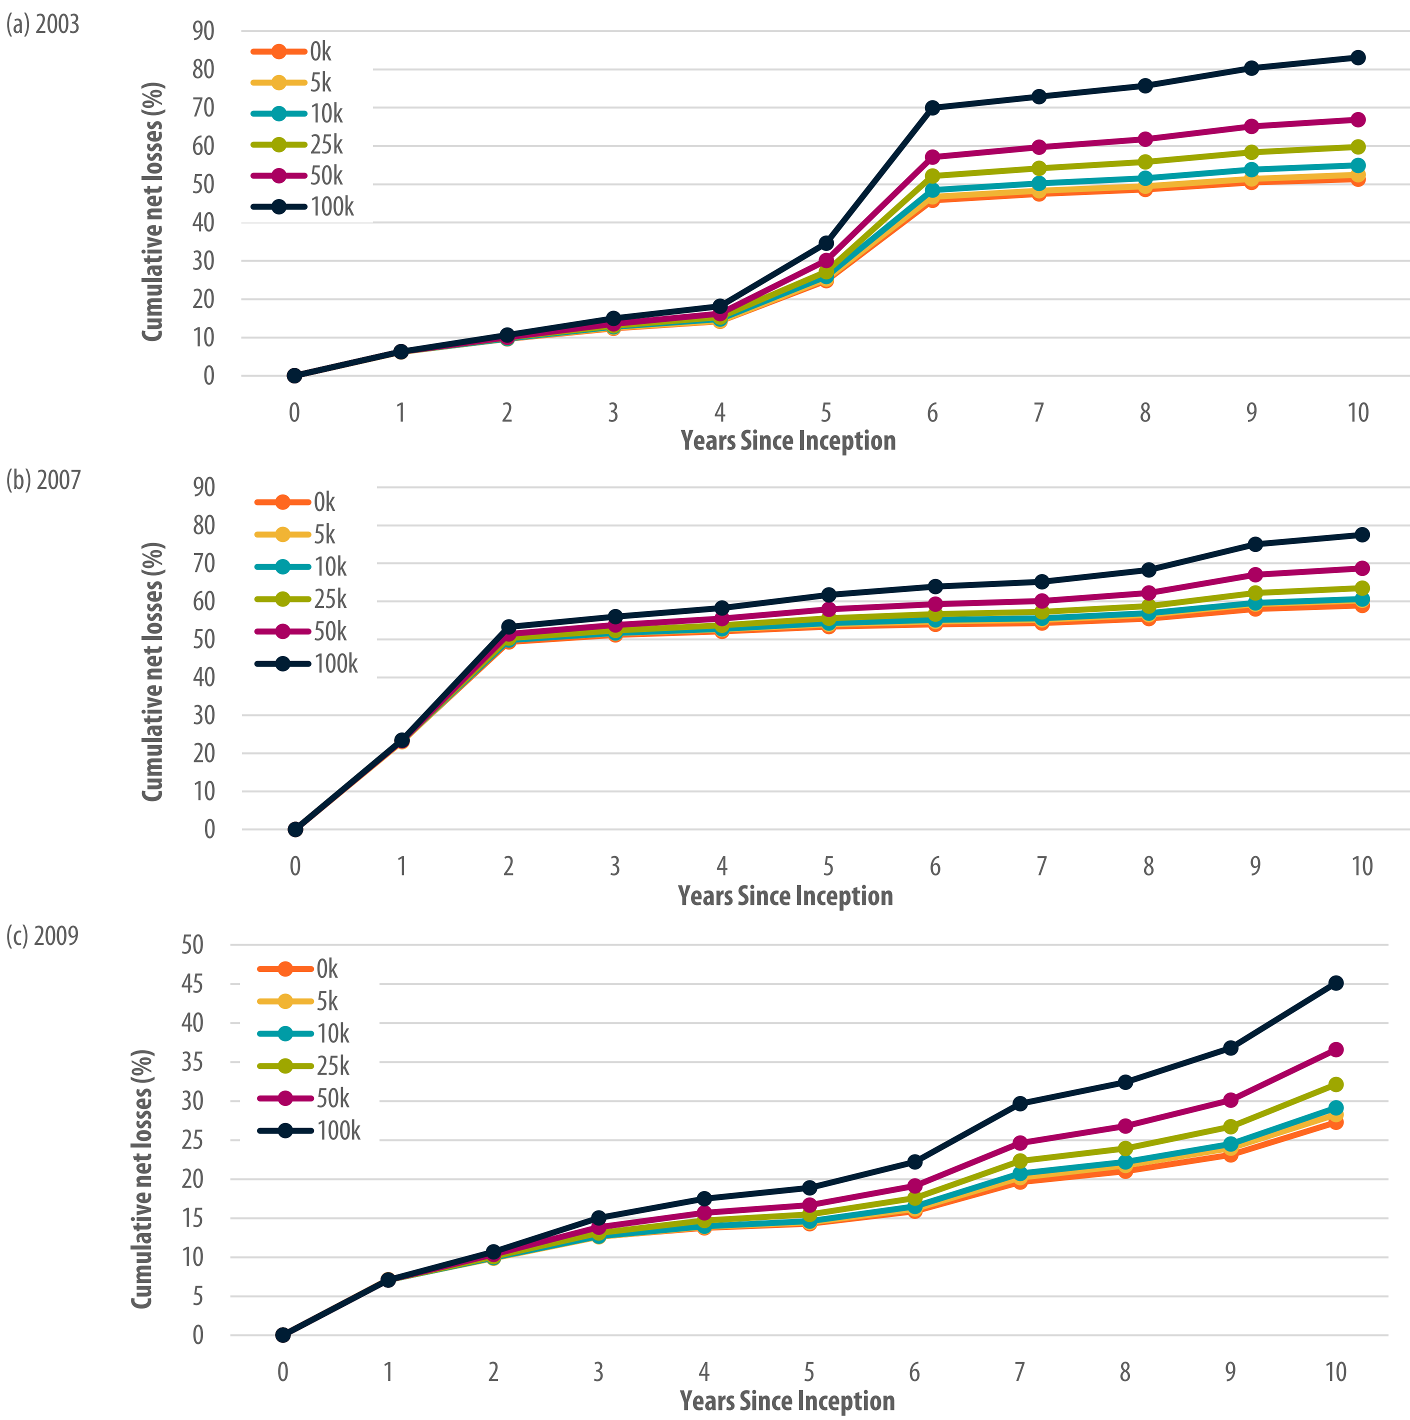

Using the scenario labels we defined before, these charts correspond respectively to *bear market occurring later*, *bear market occurring earlier*, and *bull market throughout*. The analysis given earlier about how a DI account compares in TLH to an ETF with the same holdings still holds in these scenarios when the DI account receives regular cash infusions. Furthermore, accounts that experience a significant downturn show that they approach their asymptotic slant line behavior quickly after that event. However, we again observe, especially in Figure (15a), how the slopes of that asymptotic slant line appear to depend linearly on the cash infusion amount. For accounts starting in 2009, i.e., when the market remains strong throughout, we notice a continued accelerated growth, showing that the DI account has not yet approached an asymptotic slant line, no matter the level of regular cash infusions.

These backtests confirm a similar result as we found in the previous section: when the data is aggregated, possible TLH harvesting trends toward a slant line whose slope is proportional to the amount of cash infusions (with the difference between the actual curve and this asymptotic behavior decreasing as a decaying exponential). However, when disaggregating backtest bundles into similar start years, we again see how the function of TLH varies due to market conditions. Hence, it is unlikely that any individual DI account will behave like the aggregated trend.

Implications for the TLH Function General Model

We briefly comment how our results play into our proposed general model for the TLH function.

Using the TLH general model, when we use regular cash infusions, the money invested P is no longer a constant but has the form $P = p_0 + p_1 t$, where p_0 is the initial investment and p_1 is the rate of cash infusion. We can model this as a continuous function even though the infusions occur at discrete times.

The subsection entitled "TLH Function General Model" defined the concept of loss carrying capacity M, and argued that this has the form $M(t) = c(t)P$, where P is the initial investment and the relative loss carrying capacity $c(t)$ is a positive nondecreasing function whose limit c_{max} is less than 1. In theory, $c(t)$ could be a constant but it depends on market conditions, taking into account loss harvesting possibilities and other factors.

With regular cash infusions, the invested amount P is no longer a fixed amount but depends on time as $P = p_0 + p_1 t$, where p_0 is the initial investment and p_1 is the rate of cash infusions (in \$/yr). We could suspect that the loss carrying capacity is $M(t) = c(t)P(t) = c(t)(p_0 + p_1 t)$, though it is quite possible that $M(t)$ might reasonably be more complicated than this product, even for theoretical reasons. Nevertheless, our empirical observations establish that the slope of the limiting asymptote line is 0.31 times the coefficient p_1. Furthermore, our empirical analysis supports a hypothesis that the process taking the investment function $P(t)$ to loss carrying capacity $M(t)$ is linear in $P(t)$; this is a particularly nice property.

As we stated earlier, deeper analysis into the properties of the loss carrying capacity is beyond the scope of this paper. However, the latter two observations are valuable for that study.

Conclusion

Direct indexing accounts offer the investor the potential for more flexibility than mutual funds or ETFs. One of the biggest potential benefits is that tax laws allow the investor to claim tax losses on the sale of individual stocks if their value decreases below the original invested amount. This flexibility comes at a cost: the account manager must manage the account more strategically to take advantage of the possibilities that a DI paradigm can offer.

We believe account managers should rely on optimization software to rebalance the portfolio wisely. Even with using optimization software, various settings in the optimization control the tradeoff between optimizing tax credit and remaining close to the reference index the account tracks. This paper analyzed coefficients associated to this tradeoff.

Because losses can only be harvested off an initial investment, there exists a theoretical maximum of 100% of tax losses that can be harvested. From a finance perspective, 100% is absurd. Figure (8), which affirms similar results as that found in other DI research literature, appears to indicate an ossification level of 50%, but the IQR is so wide that we believe it would be irresponsible to suggest this number as something a client could expect with a high degree of confidence. Aggregating data over all years depicts a sort of grossly averaged behavior that increases up to this maximum loss capture level as a decaying exponential. However, disaggregating the data by start year shows that a far more nuanced picture tends to occur: though the general trend of decaying up toward a horizontal asymptote generally holds, market conditions affect this significantly by increasing the rate of change through increased TLH opportunities and by modifying the actual ossification level, which is also controlled by market conditions. This paper identifies three scenarios of what can occur specifically and how much better a DI account might perform over an equivalent ETF.

Regular cash infusions mitigate the ossification phenomenon. This makes sense since these infusions increase the cost basis of the portfolio. Intuitively speaking, regular infusions offer fresh cost basis in the portfolio, which allow for harvesting opportunities. Somewhat surprisingly, our backtests show that regular cash infusions not only lead to more TLH, but they also simultaneously decrease the ending tracking error. Regular infusions lubricate the DI account to give more flexibility in rebalancing to capture more TLH while decreasing the error from the tracked index. Our backtests also strongly support the hypothesis that regular cash infusions do not raise the ossification asymptote, but rather tilt it: the asymptotic TLH behavior is no longer a flat line but rather a slant line with a nonzero slope. Furthermore, our analysis also supports that the slope of this asymptotic line is proportional to the level of cash infusion at a factor of around $310/yr per $1000 of yearly cash infusion.

The analysis of strategies for DI accounts stands as an active area in current finance research. This paper offers an introduction to the topic with analysis supported by historical data. The growing science of DI account management includes many other topics. For example, analyzing turnover may lead to another facet to the management of DI strategies. We believe other topics that address an investor's unique needs, such as how to optimize a portfolio transition or the effects of values-based and ESG criteria, are situations that warrant deeper, specific study.

Appendix

This appendix explores a property of the TASR optimization software and then addresses two details about the tradeoff coefficient: how it appears in the optimizer and a minor point about definitions and implementation of tradeoff behavior in TASR.

Optimization Algorithm Behavior

Optimization algorithms often employ randomization for a seed guess. According to our optimizer documentation, "this acts as a small perturbation to the solver, and typically leads to different solution paths." Nevertheless, if there is a single optimum solution, the optimization software will always find that unique value within a numerical tolerance parameter. In practice, if two backtests with the same parameters run through TASR without timing out the optimizer, then they return identical optimal solutions. In this regard, the optimization algorithm is deterministic. Therefore, it is ideal if the analyst chooses a large enough timeout value that the optimizer does not reach this value during a bundle of backtests. However, it might not always be practical to do this, for example when a particularly small feasible region for the optimization problem renders computations prohibitively time consuming.

Experimentation with TASR reveals another surprising phenomenon: small variations in the minimum transaction size lead to unexpected deviations in TLH or tax alpha calculations. As an example, the following table shows the simulation of an account based on a US Large Cap Index, from 5/31/2002 to 5/31/2023, with all other parameters the same except varying the minimum transaction size.

Both the simulated tax losses harvested and the tax alpha over the year vary despite small changes in transaction size. In this case, the data for harvested tax losses produce an average of $78,481.94 with a standard deviation of $972.09.

Financial intuition leads us to believe that small deviations in the minimum transaction size would not affect investor behavior and therefore be economically irrelevant.

Table 1: TASR Parameter Sensitivity.

Minimum Transaction	Taxes Offset	Tax Alpha
$99.50	$78,549.00	8.57%
$99.60	$80,385.81	8.87%
$99.70	$78,840.32	8.65%
$99.80	$77,828.26	8.56%
$99.90	$78,281.81	8.51%
$100.00	$78,260.28	8.51%
$100.10	$77,304.03	8.47%
$100.20	$78,972.58	8.60%
$100.30	$79,992.68	8.75%
$100.40	$77,281.66	8.42%
$100.50	$77,604.93	8.43%

By shrinking the step size even further, we do not observe any decrease in the standard deviations around the averages. Table 2 depicts averages and relative standard deviations of tax losses harvested and tax alpha for a strategy that tracks a US Large Cap Index for the three years (2008, 2012, and 2022, running from 5/31 to 5/31 the next year), where we ran minimum transaction sizes of 99.5 / 99.6 / . . . /100.5 (step sizes of 0.1 around 100) versus 100 / 100.01 / . . . / 100.09 (step sizes of 0.01 around 100).

Table 2: Comparison of Minimum Transaction Step Sizes.

		2008		2012		2022	
	Step size	**Average**	**Relative σ**	**Average**	**Relative σ**	**Average**	**Relative σ**
Tax loss harvested	0.1	$193,232.13	2.08%	$19,515.77	0.51%	$78,481.94	1.24%
	0.01	$190,904.99	2.54%	$19,569.73	0.48%	$78,011.96	1.22%
Tax alpha	0.1	17.10%	2.11%	2.16%	0.46%	8.58%	1.52%
	0.01	16.94%	2.31%	2.16%	0.42%	8.53%	1.57%

In all six instances, comparing the 0.1 step size to the 0.01 step size, we cannot reject with 99% confidence that the averages represent different populations. In other words, the estimated averages and the standard deviations of the estimates are not significantly different when we shrink the step size.

A geometric perspective on how the minimum transaction size affects the feasible region for the optimization might lead us to surmise that the variability (measured by relative σ) in the output data corresponding with a given step size is affected significantly by the number of securities in the portfolio. Table 3 provides data for similar time frames, with minimum transaction sizes of 99.5 / 99.6 / . . . / 100.5, but tracking a prominent index that is comprised of only 30 securities.

Table 3. Persistence of Step Size Effect in a Portfolio with Fewer Securities.

		2008		2012		2022	
	Step size	**Average**	**Relative σ**	**Average**	**Relative σ**	**Average**	**Relative σ**
TLH	0.1	$166,068.06	1.179%	$12,631.79	0.177%	$52,227.56	1.669%
Tax alpha	0.1	16.28%	1.013%	1.58%	0.182%	6.04%	1.598%

The variability in the results (measured by relative σ) due to working with a smaller number of securities in a portfolio decreases by a half for 2008 and 2012 but increases in 2022.

The data suggests the following conclusions about the variability in the output results (in particular, tax losses harvested and tax alpha) in relation to small changes in the minimum transaction size:

- remains unaffected by step size (assuming the step size is small enough);
- is not closely correlated with the number of securities in the index;
- appears more closely correlated to market conditions, in particular volatility;
- and (though we have not produced data to analyze this) is likely related to the length of time of the backtest.

Though the backtest algorithm is deterministic (assuming no optimization times out), the results appear to depend sensitively on some initial conditions (e.g., minimum transaction size). The variability of this sensitivity does not appear closely correlated with the number of securities in the portfolio, but instead appears correlated more with market conditions, in particular volatility.

As an upshot of these observations, we can take advantage of the sensitive dependence on initial conditions by varying the minimum transaction size slightly to generate backtests that capture variability inherent in market movements.

Transaction Costs Analysis

TASR's objective function incorporates the transaction costs for rebalancing a DI portfolio. Higher transaction costs ostensibly could alter the behavior of the optimization selections and modify the analyses developed in this paper. To study this issue, we used TASR to backtest a DI account tracking an SPX index, but letting the transaction costs range from 1 to 50 bps. It is common for transaction costs to lie below 10 bps.

Figure 18. SPX Index Last 5 Years

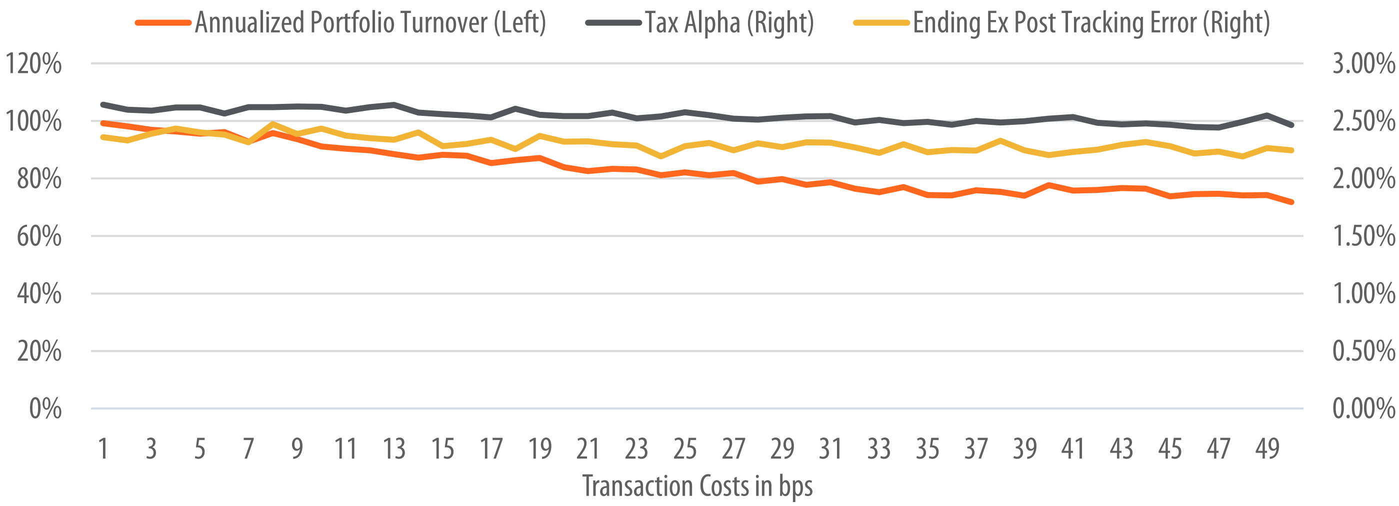

As we might expect, Figure (18) does indicate a decline in the annualized portfolio turnover. However, both the tax alpha and the ending ex-post tracking error are not significantly affected by the transaction costs. These measures for tax-aware performance affirm that the transaction costs, especially costs up to 10 bps, do not substantively change the analyses in this paper.

Pretax Performance

One concern an investor might have regarding tax alpha is whether it might be outweighed by inferior pre-tax returns. We considered the pre-tax return differences of the portfolio versus the shadow portfolio across all backtests in the long-term effectiveness study and found a median of -11 bps with 25th and 75th percentile values of -24 bps and +2 bps, respectively. The after-tax return of the portfolio, assuming top marginal tax rates, was a median of 1.4% above the after-tax return of the shadow, with 25th and 75th percentiles of 1 and 1.9%. While the slightly negative pre-tax performance is undesirable, it is significantly outweighed by the after-tax benefits achieved by smart tax management.

Tradeoff Coefficient Analysis

Two parameters named Tradeoff Value and Tradeoff Location serve different accounting purposes. However, their relationship with the tradeoff coefficient is

$$tradeoff\ coefficient = 2 \times (tradeoff\ value) \times (tradeoff\ location).$$

To run a backtest with a specific tradeoff coefficient value, any combination of Tradeoff Value and Tradeoff Location can be chosen. In this study, we kept Tradeoff Location fixed at 0.03 but varied Tradeoff Value.