Journal Of Investment Management, Vol. 14, No. 1, (2016), pp. 26–64


www.joim.com

# TAX-COGNIZANT PORTFOLIO ANALYSIS: A METHODOLOGY FOR MAXIMIZING AFTER-TAX WEALTH

*Kenneth A. Blay*[a] *and Harry M. Markowitz*[b]

*The most prevalent methods of incorporating taxes into the portfolio construction process are the preliminary adjustment of asset allocation inputs for taxes and the post-optimization application of asset location heuristics. We argue that these methods are unsatisfactory in that they fail to address taxation dynamics that result from investment and consumption-dependent illiquidities. Tax-Cognizant Portfolio Analysis (TCPA) is proposed as a methodology that addresses these issues while seeking to maximize expected after-tax wealth for given levels of risk. TCPA achieves this through the use of simulation methods to assess the impact of portfolio turnover, sequence of investment returns, and wealth consumption decisions on after-tax wealth outcomes from taxable, tax-deferred, and tax-exempt accounts.*

## 1 Introduction

Taxes are one of the most persistent frictional costs investors face. Unfortunately, asset allocation methods used by investment practitioners have generally been implemented with either complete or varying degrees of indifference towards the impact of taxes on long-term wealth. The most prevalent approaches to incorporating taxes into the portfolio construction process in use today are the preliminary adjustment of asset allocation inputs for taxes and the post-optimization application of asset location heuristics. However, the taxation of investment returns over time is a dynamic process that introduces complexities that these approaches do not fully address.

Aside from accounting for an investor's specific rates of taxation, it is also necessary to consider the taxation characteristics of the accounts in which investments are held (i.e., taxable, tax-deferred, or tax-exempt), the composition of the returns generated by investments (i.e., income, dividends, and/or capital appreciation), the sequence of returns experienced by investments, and the timing of taxation events

[a]State Street Associates, Vice President, Investment and Risk Research, 140 Mount Auburn St., Cambridge, MA 02138, USA. Email: KABlay@statestreet.com. This work was completed while the author served as Director of Research for 1st Global, Dallas, TX.
[b]Harry Markowitz Company, 1010 Turquoise Street, San Diego, CA 92109, USA. Tel. (858)488-7212 (office).
E-mail: HarryHMM@aol.com

(e.g., the realization of gains or losses that result from portfolio turnover or withdrawals from tax-deferred accounts). The occurrence of taxation events is driven by an investor's asset management and wealth consumption decisions. These decisions introduce the complexities of asset illiquidity to the portfolio analysis problem.

How assets are managed over time can have a significant impact on after-tax wealth outcomes. High turnover investment strategies expose investors to more frequent taxation than what might be experienced with low turnover strategies. Furthermore, taxable gains are not the only outcomes that result from the partial or complete sale of investments over time. Losses may be incurred that can serve to offset future realized gains. These losses and the benefits of the tax offsets they can provide must be considered to determine the true impact of taxation from portfolio turnover.

How wealth is consumed also affects after-tax wealth outcomes. Consumption from tax-deferred and tax-exempt accounts is limited by penalties that apply until an investor reaches a specific age. This means that consumption may not begin until some point far in the future. Additionally, the assumption that an investor will consume the entirety of their wealth at a single point in the future is unrealistic for determining the impact of taxation on wealth outcomes. A more realistic assumption is that an investor will consume wealth periodically, over a number of years.

The inclusion of asset management and wealth consumption variables significantly alters how an investor must approach the portfolio analysis process when seeking to maximize after-tax wealth. It also changes the objective of a portfolio analysis from seeking to maximize end of period wealth to that of maximizing the cumulative value of the after-tax cash flows that an investment can provide over an investor's lifetime based on specific investing and consumption decisions.

In this paper we introduce Tax-Cognizant Portfolio Analysis (TCPA) as a methodology that seeks to maximize after-tax wealth by addressing taxation dynamics that result from investment and consumption-dependent illiquidities. TCPA incorporates taxation and illiquidities into the portfolio analysis process by seeking to maximize the expected present value of future after-tax consumption cash flows for given levels of risk. Expected present values are determined through a series of concurrent simulations of an investor consuming accumulated wealth from asset class investments as if they were held in each of the possible account types included in the analysis. The means, standard deviations, and correlations of simulated present values (PVs) are then used as inputs for a mean–variance portfolio analysis. The process simultaneously identifies the approximately optimal allocations to asset class investments and the location of those investments across account types with different taxation characteristics.

Our introduction of TCPA begins in Section 1 where we discuss the current state of tax-aware asset allocation and how TCPA differs from current approaches. In Section 2 we detail the derivation of the inputs needed to develop a tax-cognizant efficient frontier. Here we describe the necessary specifications for three component models (lifecycle, consumption, and taxation) that are required for simulating the present values of after-tax consumption cash flows. We also address the issue of estimation error in determining means, standard deviations, and correlations of present values from a series of simulations.

In Section 3 we produce a PV mean–PV standard deviation tax-cognizant efficient frontier. Here we begin to demonstrate the benefits and characteristics of tax-cognizant efficiency by comparing the present value frontier with a frontier developed using mean–variance optimization with standard inputs that have not been adjusted for taxes. Asset location preferences exhibited by the TCPA process are reviewed. We also address issues with instability in the composition of TCPA frontiers that result from simulations producing asset class present value outcomes that are close substitutes. We conclude Section 3 by identifying a critical impracticality of the present value frontier: PV mean and PV standard deviation are not intuitive parameters for portfolio selection.

In Section 4 we address the impracticality of the PV frontier by transforming it into more intuitive terms. We work through the conversion of expected present values to expected average annual real after-tax cash flows. Rather than convert expected PV standard deviations we detail two transformations of the PV frontier through the use of left-tail confidence levels. We close Section 4 by presenting two distinct approaches to tax-cognizant portfolio selection that allow investors to select portfolios based on the trade-off between average annual real after-tax cash flow and the likelihood of achieving the cash flow.

In Section 5 we discuss the time-dependent nature of TCPA and implementations of tax-cognizant investing that balance tax efficiency with the risk of not achieving desired after-tax wealth outcomes. Two dynamic approaches are presented: a constant confidence level strategy and an increasing confidence level strategy. We present our concluding remarks in Section 6.

## 2 The current state of tax-aware asset allocation

With the persistence and significant impact of taxation on the amount of wealth an investor can accumulate and ultimately consume, it is surprising that the literature on tax-aware asset allocation is not more complete. Fortunately, tax-aware investment management has been receiving increasing attention. Horan and Adler (2009) and Jennings *et al.* (2011) provide a thorough review of the existing literature on tax-aware investment management and its implementation in practice. More recently, Blanchett and Kaplan (2013) identified tax-efficient decisions, including asset location, as a fundamental part of financial planning that can have a significant impact on retirement income. While the importance of minimizing the impact of taxation on wealth outcomes has been gaining recognition, the adoption of tax-aware asset allocation methods by practitioners generally remains limited to the preliminary adjustment of asset allocation inputs for taxes and the post-optimization application of asset location heuristics.

The most prevalent methods for tax-adjusting optimization inputs are those presented by Reichenstein (2004), Wilcox *et al.* (2006), Horan (2007), and more comprehensively by Horan and Al Zaman (2008). Although the notion of adjusting inputs for expected taxes might seem like a straightforward proposition, this is far from the case. Each of the methods from the authors above approaches the adjustments to expected returns, standard deviations, and correlations differently. Adjustments are also dependent on whether a portfolio analysis is being conducted using the pre-tax or after-tax values of an investor's current assets.

Reichenstein and Horan recommend determining asset allocations using after-tax values. An example of after-tax valuation is the treatment of tax-deferred accounts. Rather than assuming that an investor with $100,000 in a traditional Individual Retirement Account (IRA) can use $100,000 for future consumption, one should reduce the value of the account by the expected (embedded) tax liability that will result when assets are consumed. If the expected marginal income tax rate in retirement is assumed to be 25 percent, the after-tax value of the IRA account is calculated to be $\$100,000 \times (1 - 0.25)$, or \$75,000. While there is generally agreement regarding the valuation of tax-deferred and tax-exempt accounts, there is disagreement regarding the valuation of taxable assets. Reichenstein (2000, 2001, 2006, 2007a, 2007b) believes that no tax adjustment is necessary for taxable accounts while Horan (2007) believes that an adjustment is appropriate.

An important concept to note when adjustments are made to asset allocation inputs is that the same investment held in three different types of accounts (i.e., taxable, tax-deferred, or tax-exempt) should be treated as three distinct investments. Tax adjustments to tax-deferred accounts are made only if the analysis is being conducted on a pre-tax basis. In such a case, a tax liability exists that requires adjustments to expected returns, standard deviations, and correlations. When the analysis is conducted on an after-tax basis no adjustments to inputs are necessary and tax-deferred assets are treated as if they were tax-exempt since taxes have already been addressed.

Adjustments to expected returns are generally accomplished by subtracting the expected tax liability for asset class investments when held in each of the account types considered. The tax liability used can vary across methods based on the variables used to determine the effective tax rate, the type of investor that is assumed for the analysis (i.e., trader, active investor, or passive investor), the specific expectations for realized capital gains, or whether taxes are subtracted from the entire pre-tax return or the portion of the return that represents a risk premium. Adjustments to standard deviations are generally based on the tax rate used to adjust returns. Adjustments to the correlation or covariance matrix, not surprisingly, are closely related to the changes made to standard deviations. Interestingly, Horan changes the covariance matrix by applying a different method than Wilcox, Horvitz, and diBartolomeo while Reichenstein makes no changes.

It is evident from reviewing the literature on the tax-adjustment of asset allocation inputs that there is agreement that taxes affect how the portfolio optimization problem is approached and that there are significant benefits to considering taxes in portfolio construction. Unfortunately, there are important areas where disagreement persists. There is also the problem that neither of the methods presented by the previously identified authors directly address the complexities of illiquidity introduced by an investor's asset management and consumption decisions.

Markowitz (1959) understood that incorporating taxes into a portfolio analysis was not as simple as merely adjusting inputs for the effects of taxes. He explained that *"taxes would pose no special problem of analysis if it were not for the provision that capital gains are taxed only when 'realized' (i.e., only when the security is sold). Without this consideration we could simply define means, variances and covariances in terms of return after taxes, and perform an ordinary efficient set analysis with our inputs thus defined."* Problems with illiquidity arise from the fact that after-tax wealth outcomes depend not only on the return achieved by an investment but also on the negative impact of taxes on an investment's return. Taxable events, such as the partial sale of an investment due to portfolio turnover or investor wealth consumption, may occur regularly or randomly. In either case, it is uncertain whether the investor will experience a gain or a loss. The timing of the taxation event and the determination of whether an investor experienced a gain that is taxed, a gain that is offset by previous losses, or a loss to serve as an offset for future taxable gains is dependent on the sequence of returns experienced by the investment up to the time that the taxable event occurs.

As a possible method of dealing with illiquidity issues Markowitz (1959) proposed defining the "worth" of a security at the end of a period by assessing its market value and then subtracting taxes payable on realized gains and income. Additional adjustments are then made to reflect "likely" taxes on existing capital gains and "likely" savings from existing capital losses in the future. The extent of these adjustments could be based on past experience regarding gains and losses or judgment regarding the particular characteristics of an investment. An analysis would then be conducted using the means, variances, and covariances of the worth of securities defined in this manner. Markowitz presented this proposal well before the formalization of 401(k) plans and the introduction of tax-deferred and tax-exempt IRAs. Assets held in these types of accounts generally limit withdrawals until a specified retirement age. These limitations represent additional consumption illiquidities.

The second most common approach to tax-aware asset allocation represents a combination of asset allocation with asset location methods. Asset location methods seek to provide investors with guidance as to where to hold specific investments across different account types so as to achieve the greatest after-tax benefit. However, the asset location decision is often implemented using general rule-of-thumb strategies **after** the asset allocation decision has been made based on inputs that have not been adjusted for taxes. That is, expected after-tax investment risk, returns, and relationships were not considered as part of the optimization process. In effect, this approach is the equivalent of changing inputs and expecting there to be no change in the composition of the efficient set of portfolios. Consequently, the approach results in suboptimal portfolios.

Shoven and Sialm (2003) and Damon *et al.* (2004) have advanced the understanding of asset location by providing an analytical framework for making asset location decisions. In practice, however, there continues to be disagreement as to which assets are best held in specific accounts. Horan and Adler (2009) present the results of a 2008 survey of CFA Institute members who manage taxable accounts or were familiar with how their firm managed taxable accounts. The survey showed that while most investment managers preferred to place taxable bonds in tax-deferred accounts, more than a quarter preferred placing taxable bonds in taxable accounts. Despite an improved understanding of the benefits of asset location, the post-optimization application of subjective asset location heuristics results in suboptimal portfolios. The approach also fails to consider investment and consumption-dependent illiquidities.

The TCPA methodology we present in this paper advances the current state of tax-efficient asset allocation in that it directly addresses taxation dynamics that result from investment and consumption-dependent illiquidities. Additionally, the approach does not require the determination of an investor's true after-tax asset allocation or the application of generic asset location heuristics.

## 3 Deriving tax-cognizant optimization inputs

The first step in the tax-cognizant portfolio selection problem is establishing expectations about the after-tax cash flows provided by asset class investments. Tax-exempted[1] inputs are an investment practitioner's expectations for the return, risk, and correlation characteristics of investments included in a portfolio analysis. To translate these inputs into expectations about after-tax cash flows, and ultimately tax-cognizant optimization inputs, it is necessary to have a process for determining the impact of taxes on the growth and consumption of asset class investments.

The inputs for a TCPA are a result of the interaction between tax-exempted inputs and three component models. Tax-exempted inputs are used as parameters in determining the returns provided by asset class investments over an investor's simulated lifetime. Simulated lifetimes are based on an investor lifecycle model that incorporates a wealth consumption model and a taxation model. The wealth consumption model determines the amount of the cash flows provided by asset class investments. The taxation model determines the impact of taxes on investments and cash flows. The present value of the after-tax cash flows for each simulation iteration is then determined. The means, standard deviations, and correlations of the present values of after-tax consumption cash flows from a series of simulations ultimately become the inputs to a TCPA. In this section, we delineate the elements necessary to develop tax-cognizant inputs and detail the required specifications for the process. As we progress through the TCPA process we provide the information used to develop the tax-cognizant efficient frontiers presented later in this paper.

### *3.1 Tax-exempted asset allocation inputs*

The starting point for a TCPA is the development of tax-exempted asset allocation inputs. Specifically, a standard portfolio analysis requires forward-looking assumptions for means, standard deviations, and correlations for all asset classes that are included in the analysis. The methods used to develop these inputs may be amongst any of those traditionally employed by competent asset allocation practitioners. However, it is important that tax-adjustments are not made to inputs as they are used as parameters for a Monte Carlo process to simulate asset class returns which are subsequently taxed using a taxation model.

### *3.2 Account types*

A fundamental aspect of a TCPA is the identification of the account types included in the analysis and the specification of the taxation characteristics of those account types. Our analysis will include the three most common types of accounts available to investors. The accounts and their taxation characteristics are as follows:

**1. Taxable:** These accounts are funded with after-tax dollars. Investment returns are subject to both capital gain and income taxes.
**2. Tax-deferred:** These accounts are generally funded with pre-tax dollars. Investment returns are not subject to taxation. However, withdrawals from these accounts are subject to taxation as income at the rate that is applicable to the investor at the time of withdrawal. Examples of this type of account include traditional IRAs and 401(k) accounts.
**3. Tax-exempt:** These accounts are funded with after-tax dollars. Investment returns and withdrawals are generally not subject to taxation. Examples of this type of account include Roth IRA and Roth 401(k) accounts.

Along with specific taxation characteristics, each of these accounts also has unique liquidity characteristics. Taxable accounts are assumed to be liquid because withdrawals can generally be made at any time without penalty. However, the fact that gains are subject to taxation imposes liquidity limitations. In effect, taxes can be viewed as a penalty for selling or exchanging investments. Assets held in tax-deferred and tax-exempt accounts are not subject to the illiquidity imposed by taxation but they do have age-dependent liquidity limitations. Withdrawals from these accounts may be subject to both taxation and penalties if they occur before an investor reaches the age of 59½. While there are exceptions to these withdrawal penalties, we assume that an investor will only begin withdrawing from these accounts at a retirement date which occurs after penalties no longer apply. The liquidity limitations presented by each of these accounts are directly addressed by one or a combination of the component models used to derive TCPA inputs.

### 3.3 Expected tax rates

The most important variables in determining the impact of taxes on the cash flows provided by asset class investments are the tax rates that will apply to those investments. However, future tax rates are not known with certainty. Given that an investor can only act on expectations of future tax rates, it is reasonable to proceed with the assumption that existing rates will persist into the foreseeable future with the understanding that another analysis must be conducted if rates change.

While future tax rates are uncertain, there are changes in our investor's tax circumstances that can be anticipated. It is possible that an investor's applicable tax rates will be higher when they are working and accumulating retirement assets than when they are consuming assets in retirement. The lifecycle model used in determining after-tax cash flows divides an investor's lifetime into an accumulation period and a consumption period. This allows for the specification of different tax rates for each of these time periods. Appropriate state and local taxes should be aggregated with federal tax rates to determine the applicable rates. Tax rate specifications should include the income tax rates ($T_I$) that apply to ordinary income and qualified dividends, short-term capital gain tax rates ($T_S$), and long-term capital gain tax rates ($T_L$). For our example, we assume that our investor has a 33 percent marginal tax rate during accumulation and a 28 percent rate in consumption. Tax rates for income and capital gains are specified in more detail later.

### 3.4 Composition of asset class returns

Determining how taxes impact returns requires the delineation of the components of asset class expected returns into return from capital appreciation and return from income. This is because capital gains and income can be taxed at different rates. Additionally, income is normally a steady recurring return while capital gains are more volatile and require the sale of an investment to be realized. These are important characteristics that must be considered when simulating after-tax cash flows. Because the asset class return assumptions that are used as simulation parameters are based on total returns, it is necessary to decompose the total return for each asset class into its component parts. This is accomplished by specifying the average percentage of an investment's return that is derived from capital appreciation and the percentage derived from income. Historical data can be used as a basis for estimation of these component percentages.

Variability in the amount of income produced by an asset class is to be expected over time. However, it should be representative of the income normally produced by the asset class. In order to accommodate for this variability and respect the attributes of how an asset class generates income, simulations for income returns should be conducted concurrently with simulations for total returns. This can be done using the expected income return estimate (determined as the income portion of the expected total return) and estimates for the variability of income. The relationship between the income and capital appreciation of an asset class should also be respected. This can be approximated by specifying the correlation between income and total return as a simulation parameter. Again, historical data can serve to inform forward-looking estimates for income volatility and correlation with total return. Subtracting income returns from total returns allows for the determination of the returns from capital appreciation in simulation runs.[2] This approach produces simulations of each of the components of an asset class investment's returns in a manner that conforms to the initial tax-exempted inputs.

### 3.5 *Asset class taxation characteristics*

The taxation of an investment's returns depends on the taxation characteristics of the component parts of the returns. Asset classes can generate different types of income that are taxed at different rates. As such, the type of income generated by asset classes should be specified as ordinary taxable income, tax-exempt income, or qualified dividends so that taxes are applied accordingly.

Another important consideration is an investor's preference for implementing asset class investments. If an investor implements using passive index investments, turnover expectations are usually low. If an investor implements using actively managed investments, turnover expectations will generally be higher. Although turnover is only pertinent to assets held in taxable accounts, it is an important component in determining how taxes affect investment outcomes. Both short-term and long-term turnover expectations are specified for every asset class to account for the notable differences in the tax impact between the two. Short-term turnover ($TO_S$) is the percentage of an investment that is expected to be held less than one year. Long-term turnover ($TO_L$) is the percentage of an investment held for longer than one year that is expected to be sold in any given year. Our research indicates that most of the negative tax impact of turnover occurs at levels as low as 30 to 40 percent. Once these levels of turnover have been exceeded, the majority of the damage to investor wealth that results from taxation has been done. These results coincide with Jeffrey and Arnott (1993).

### 3.6 *The investor lifecycle model*

The investor lifecycle model is a central component in determining the appropriate inputs for a TCPA. This is because it describes how an investor is expected to act while growing and then consuming wealth. The lifecycle model incorporates a wealth consumption model that details how an investor will consume assets and a taxation model that approximates how taxes impact an investment's growth and cash flows. Inputs derived from the use of these models are a direct result of the specific parameters incorporated and thus are pertinent to an investor only if the lifecycle model used is representative of their expected tax circumstances, investment management preferences, and wealth consumption behavior.

The lifecycle model presented describes how most investors approach planning for and living through retirement. The model separates an investor's lifetime into two distinct time periods: an accumulation period and a consumption period. During the accumulation period an investor's objective is to grow and accumulate wealth in preparation for future consumption. During the consumption period an investor continues with the objective of growing wealth but now periodically consumes wealth. Figure 1 presents an example of the lifecycle model using a static rate of return. Our analysis assumes that our investor is 30 years from retirement and expects to live 30 years in retirement.

### 3.7 *Wealth consumption*

After-tax wealth for a TCPA is defined as the cumulative value of the after-tax cash flows that an investment can provide over an investor's lifetime. How the investor decides to consume their wealth over time is a key determinant of the

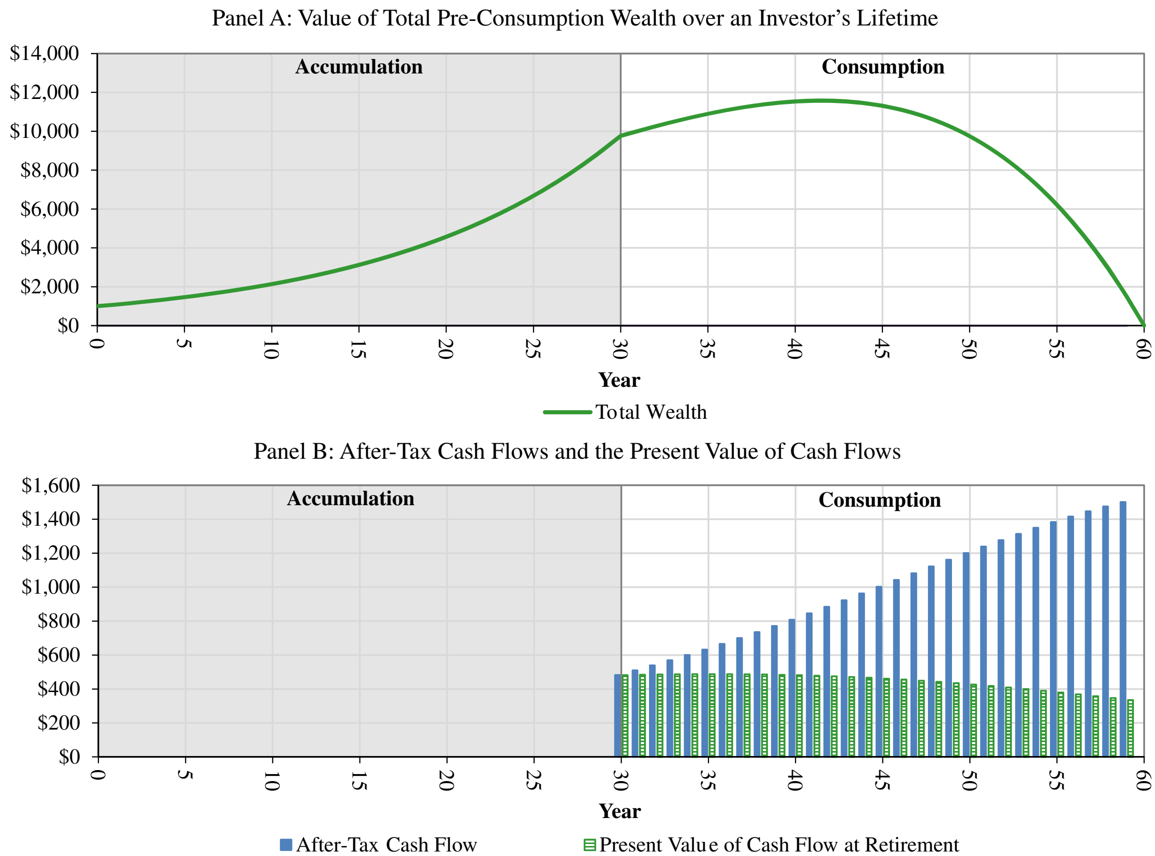


**Figure 1** An example of the investor lifecycle model assuming a static rate of return, 30 years in accumulation and 30 years in consumption. Panel A shows the value of $1,000 invested in an asset class investment over the investor's lifecycle as it grows and is then consumed. Panel B presents the after-tax consumption cash flows provided by the asset class investment and the present values of the individual cash flows discounted to the beginning of consumption period.

amount of after-tax wealth an investment can be expected to provide. To evaluate after-tax cash flow outcomes on an equivalent basis it is necessary that the wealth consumption method used by the investor be applied consistently across all asset classes and account types. Identifying portfolios of asset class investments that maximize expected after-tax cash flows requires that consumption from individual asset class investments translate directly to consumption from portfolios of asset class investments.

To meet this requirement, we have applied a modified fractional wealth consumption model in our analysis. In this model, the determination of the fraction of wealth consumed by the investor in any given year begins by calculating simple fractional consumption which is equal to one divided by the number of years remaining in the consumption period. For example, if the consumption period has been determined to be 30 years in length then the fraction of wealth consumed in the first year using simple fractional consumption will be 1/30, in the second year 1/29, in the third year 1/28 . . . and so on.

While simple fractional consumption has the benefit of being easy to understand and calculate, it provides for distribution cash flows that in some instances fail to meet Required Minimum Distributions (RMDs) from traditional IRAs as imposed by the United States Internal Revenue Service.[3] In order to avoid the possibility of onerous penalties of 50 percent on the difference between actual and required distributions, it is necessary to modify consumption so that RMDs are achieved.

This is accomplished through the modification of simple fractional consumption using a forward consumption rate ($F$) and a forward consumption dampening rate ($D$).[4] The forward consumption rate is the rate at which an investor pushes possible future consumption forward into the current consumption period. Continuing from the example above, the fraction of wealth consumed in the first year is now modified to be $1/30 \times (1+F)^{30-1}$, in the second year $1/29 \times (1+F)^{29-1} \ldots$ and so on. A possible consequence of applying a forward consumption rate is that an investor could consume a significant portion of their assets early in the consumption period. Therefore, it is necessary to limit the impact of forward consumption by also including a consumption dampening rate in the modification.

The forward consumption dampening rate is the rate at which $F$ is reduced to limit overconsumption early in the consumption period. Once again continuing from the previous example, the fraction of wealth consumed in the first year is now further modified to be $1/30 \times [1 + (F/(1 + D)^{30-1})]^{30-1}$, in the second year $1/29 \times [1 + (F/(1 + D)^{29-1})]^{29-1} \ldots$ and so on. $F$ and $D$ can be modified as necessary to meet various consumption preferences. For our analysis we have set $F$ to 3 percent and $D$ to 2.75 percent to conform with single life expectancy RMDs.

Fractional consumption modified in this manner provides for consumption cash flows that achieve RMDs for a wide range of investors who begin their consumption no earlier than age 60 and plan for a consumption period that extends no longer than age 105. It also avoids having to incorporate a specific retirement age into the analysis.

A review of Figure 1, which applies the modified fractional consumption model as implemented in our analysis, shows steadily increasing cash flows over time when a constant rate of return is assumed. This increase can serve to offset the impact of inflation. When returns are random, the model seeks to balance prudent wealth consumption and compliance with RMDs. This approach is provided as a reasonable example of a method of consuming wealth. An alternative consumption approach may be incorporated as long as it meets the necessary condition that consumption from individual asset class investments translates directly to consumption from portfolios of asset class investments.

An implication of using modified fractional consumption as specified is that all of the investor's assets will have been consumed by the end of their expected lifetime. Incorporating this assumption allows us to maximize the entirety of our investor's after-tax wealth, but it also presents a fundamental risk in wealth consumption that should be addressed.

For our analysis we specified that the investor expects to live 30 years in retirement. If we assume that the investor's age at retirement is 65 then this analysis expects consumption to continue until age 95. This consumption period was selected solely for illustrative purposes. A TCPA practitioner might consider adopting a more conservative approach of planning for longer consumption periods than indicated by our example, given the implications of an investor outliving the whole of his/her investable wealth.

### *3.8 The Blay–Markowitz taxation model*

The process of simulating after-tax cash flows from investments held in tax-deferred and tax-exempt accounts is straightforward. Income and gains have no tax implications for these accounts. Moreover, taxes are only a consideration for tax-deferred accounts when assets are withdrawn. The tax liability in such cases is calculated as the amount of the withdrawal multiplied by the investors expected marginal income tax rate.

Taxable accounts introduce various complexities. We have identified a number of these in the process of determining the taxation elements of a TCPA including the specification of appropriate tax rates that apply to an asset class investment's income, short-term capital gains, long-term capital gains, and the determination of the rates of portfolio turnover that generate capital gains. Furthermore, taxation in taxable accounts is a dynamic process where taxes are dependent not only on the returns achieved in any given year but also on the returns experienced in preceding years. More specifically, it is not enough to consider the capital gains realized in one year without considering the possibility of the tax-offsetting benefits of losses realized and carried forward from previous years. To account for these complexities, it was necessary to develop a taxation model.

The Blay–Markowitz taxation model was developed as a general model for the taxation of asset class returns over time. It is intended to provide a reasonable approximation of the impact of taxes on an investor's ability to grow and then consume wealth from an investment over an investor's lifetime. The model is much simpler than the tax code it is intended to represent in that it incorporates only the most salient elements of taxation. Specifically, the model accounts for the taxation of the income, realized short-term gains, and realized long-term gains generated by an asset class investment. Tax offsets to gains from losses carried forward ($L$) from previous years are also included in the model. However, they are only used to offset realized gains. The use of losses as an offset to income is not included as the impact of an income offset is dependent on the size of the account being considered. For example, a \$3,000 offset to income, which is the current limit on the amount of income that can be offset by losses, has a much greater impact on income generated by a \$100,000 account than when the same offset is used for a \$10 million account.

The model defines an investor's total wealth ($W_T$) as an investor's after-tax wealth ($W_A$) plus untaxed wealth ($W_U$) which represents any unrealized capital gains or losses.

$$W_T = W_A + W_U.$$

As the equation above suggests, $W_T$ is the value of an investor's total pre-consumption wealth. Wealth consumption cash flows from taxable accounts consist of equal fractions of both $W_A$ and $W_U$ assuming that a gain or loss exists in the account. That is, if 1/20th of the account is being consumed, then 1/20th of $W_A$ is consumed along with 1/20th of $W_U$. After-tax cash flows are calculated by taxing the $W_U$ portion of the cash flow at the investor's long-term capital gains tax rate. Short-term gains are addressed through the taxation of gains realized as a result of short-term turnover. Once taxed, these short-term gains become after-tax wealth ($W_A$). This method of taxing consumption is the equivalent of using the average cost method for determining the tax liability that results from the sale of a portion of an investment. Approaching the taxation of consumption in this manner avoids the need to incorporate individual tax lots into the model.

Because unrealized gains or losses are not a factor for a tax-deferred account, consumption cash flows consist of $W_T$. The values of after-tax cash flows are calculated by taxing the entire distribution at the investor's expected income tax rate during the consumption period. Distributions from tax-exempt accounts are not taxed.

Although our analysis will assume that wealth begins in an after-tax state for taxable accounts and that no loss carry-forwards exist, the model allows for taxable account simulations to begin with existing amounts of $W_A$, $W_U$, and $L$. These parameters do not affect tax-deferred accounts as distributions from these accounts are taxed at the investor's specified marginal tax rate when withdrawn regardless of whether distributions resulted from capital appreciation or not.

A more intricate taxation model that includes additional aspects of taxation can be constructed. However, a parsimonious approach is more likely to be useful to a broader range of investors. Increasing the complexity of a model may also unnecessarily increase the complexity of implementation. A detailed presentation of the Blay–Markowitz taxation model is included in the Appendix.

### 3.9 *Simulating after-tax cash flows, present values, and tax-cognizant optimization inputs*

A TCPA identifies the most beneficial account locations for asset class investments and simultaneously the most efficient allocation to those investments by first determining the after-tax cash flows that can be expected from each investment when it is held in each of the account types included in the analysis. That is to say that for the purposes of both simulation and portfolio optimization, the same asset class investment held in three different accounts should be treated as three distinct investments when each of the accounts has unique taxation characteristics. Our analysis included eight asset classes and three distinct account types. Because we have assumed that municipal bonds will only be held in a taxable account, the concurrent simulation of 22, not 24, distinct investments is required.

Rather than conduct an analysis using multiple series of future expected cash flows, it is both useful and convenient to determine how the investor values those cash flows today. Using the present values of after-tax cash flows allows us to summarize the information contained in the series of cash flows into a single number that also contains qualitative information about the investor's consumption preferences.

The interest rate used for discounting after-tax cash flows to the present is the rate that matches our investor's preference for intertemporal substitution in consumption. This is the rate at which we assume that the investor is willing to forego present consumption in order to benefit from greater future consumption. A measure for determining this consumption preference for the "typical" investor is the long-term risk-free rate. For our analysis we used (5.3 percent) the average income return of the Ibbotson Associates SBBI Long-Term Government Bond Index for the period 1926–2013.

With all of the elements necessary for simulations specified we present, in Figure 2, an example of a simulation for a large company stock fund investment held in a taxable account along with both after-tax cash flows and the present values of those cash flows.

Our analysis now advances to establishing expectations for present values, the variability of present values, and the relationships between present values. This is accomplished through the concurrent simulation of asset class investments over a series of investor lifetimes. We have provided the component models (investor lifecycle, wealth consumption, and taxation) and detailed the specifications that are necessary for

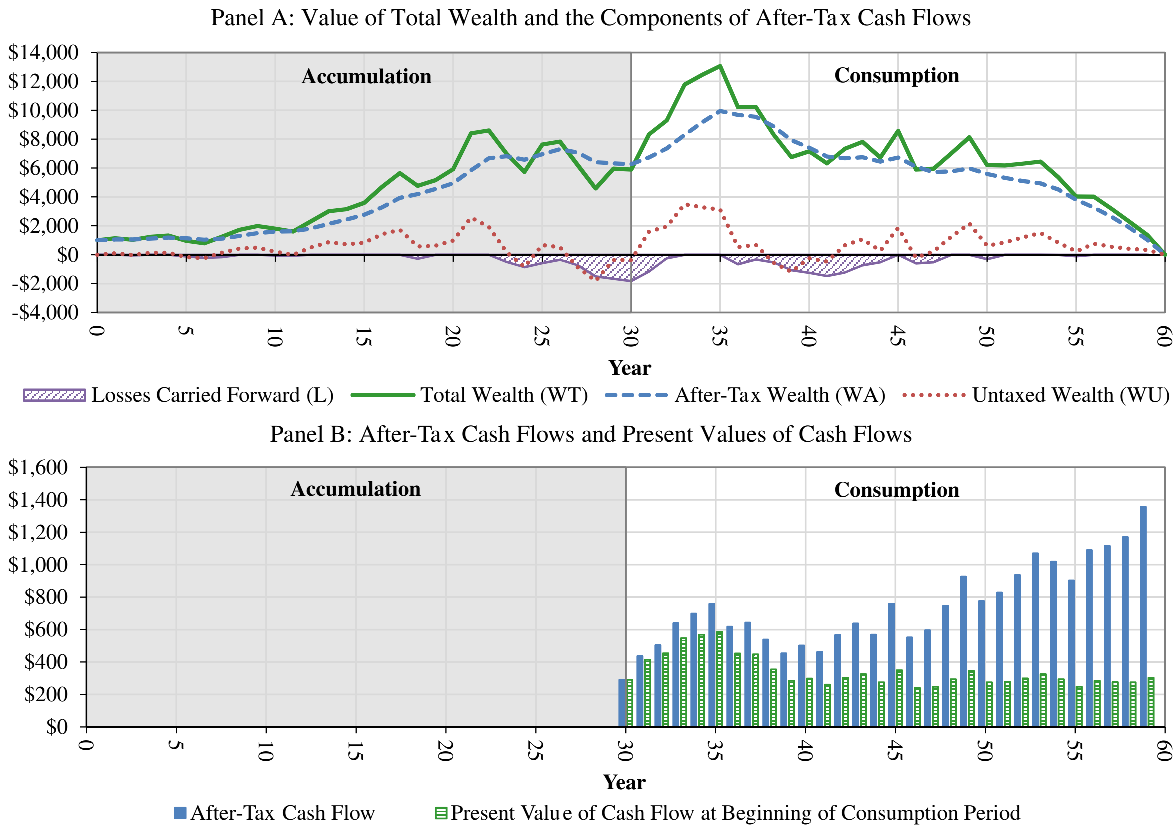


**Figure 2** An example of a lifecycle simulation for a U.S. Large Company Stock fund held in a taxable account. Panel A shows the accumulation of wealth and the state of the different components of wealth over the investor lifecycle. Panel B shows the after-tax cash flows that result as wealth is consumed along with the present value of those cash flows discounted to the beginning of the consumption period. The simulation assumes that short-term turnover ($TO_S$) = 30%, long-term turnover ($TO_L$) = 30%, the tax rate on short-term gains ($T_S$) = 33% in accumulation and 28% in consumption, the tax rate on long-term gains ($T_L$) = 15% in both periods, and the tax rate on income ($T_I$) = 15% in both periods as qualified dividends.

conducting these simulations. However, before simulations are conducted, the issue of estimation error must be addressed.

### *3.10 Estimation error*

The process of deriving TCPA inputs is not as simple as calculating the sample means, standard deviations, and correlations of the present values that are generated by the simulator. The simulation of investments compounded over multiple periods produces present value outcomes that are lognormally distributed. Unlike return outcomes, present value outcomes are bounded by zero. One of the implications of this is that outliers with large present values, particularly for high returning and highly volatile asset classes, can skew the averages of even very large samples of present values, resulting in significant sampling error. This was confirmed in attempts to achieve stable present values across multiple simulation runs. Even in simulation runs with 200,000 iterations, arithmetic averages and standard deviations of present values demonstrated notable differences from one simulation to another for the most volatile and highest returning asset classes.

**Table 1** A comparison of estimates for simulated present values, logs of present values, and present value estimates based on the logs of present values from multiple 25,000 iteration simulation runs for emerging market stocks held in a tax-exempt account.

| Emerging market stocks held in a tax-exempt account | | Simulation run (25,000 iterations per run) | | | | |
|---|---|---|---|---|---|---|
| | | 1 | 2 | 3 | 4 | 5 |
| Simulated present values | Arithmetic mean | \$10.12 | \$9.26 | \$9.62 | \$10.00 | \$9.25 |
| | Standard deviation | \$72.00 | \$38.76 | \$39.90 | \$49.41 | \$45.05 |
| Logs of simulated present values† | Arithmetic mean | 0.60 | 0.58 | 0.60 | 0.61 | 0.59 |
| | Standard deviation | 1.80 | 1.80 | 1.80 | 1.80 | 1.78 |
| | Skewness | 0.06 | 0.04 | 0.07 | 0.06 | 0.04 |
| | Excess kurtosis | 0.00 | 0.02 | 0.00 | 0.05 | 0.05 |
| Present value estimates based on logs of simulated present values | Arithmetic mean | \$9.21 | \$9.04 | \$9.21 | \$9.34 | \$8.73 |
| | Standard deviation | \$45.71 | \$44.87 | \$45.57 | \$46.44 | \$41.37 |

†Values have been rounded. Calculations used unrounded values.

A solution to the issue of estimation error can be found in the relationship between lognormal and normal distributions. A characteristic of lognormal distributions is that if a random variable $X$ is lognormally distributed then the natural logarithm of the random variable, or $\ln X$, is normally distributed. This allows us to address issues with outliers and estimation error by using the logs of simulated present values. The derivation of the necessary mean–variance inputs for a TCPA then becomes a four-step process. First, we simulate a series of present values for the asset class investments included in the analysis. Second, we take the logs of the simulated present values (lnPVs). Third, we calculate the means, standard deviations, and correlations of the lnPVs. Finally, we transform these statistics from normal terms back to lognormal terms.[5] The transformed means, standard deviations, and correlations are the TCPA inputs which are then used for conducting a mean–variance portfolio optimization.[6]

**Table 2** Summary of tax-exempted capital market expectations with asset class return and taxation characteristics.

| Asset class | | FI | MFI | LC | SC | DM | EM | RE | C |
|---|---|---|---|---|---|---|---|---|---|
| **Correlation matrix**[7] | | | | | | | | | |
| Investment grade bonds | FI | 1.00 | 0.47 | 0.22 | 0.21 | 0.14 | 0.12 | 0.17 | 0.00 |
| Municipal bonds | MFI | 0.47 | 1.00 | 0.27 | 0.25 | 0.18 | 0.16 | 0.19 | 0.00 |
| U.S. Large Company Stocks | LC | 0.22 | 0.27 | 1.00 | 0.84 | 0.75 | 0.70 | 0.62 | 0.16 |
| U.S. Small Company Stocks | SC | 0.21 | 0.25 | 0.84 | 1.00 | 0.64 | 0.67 | 0.66 | 0.18 |
| Developed market stocks | DM | 0.14 | 0.18 | 0.75 | 0.64 | 1.00 | 0.52 | 0.46 | 0.25 |
| Emerging market stocks | EM | 0.12 | 0.16 | 0.70 | 0.67 | 0.52 | 1.00 | 0.47 | 0.26 |
| Real estate | RE | 0.17 | 0.19 | 0.62 | 0.66 | 0.46 | 0.47 | 1.00 | 0.14 |
| Commodities | C | 0.00 | 0.00 | 0.16 | 0.18 | 0.25 | 0.26 | 0.14 | 1.00 |

**Table 2** (*continued*)

| Asset class | | Expected standard deviation % | *A* Expected return % | *B* % of return gain | *C* % of return income | $A \times B$ Gain return | $A \times C$ Income return | Income standard deviation % | Correl. income, total return |
|---|---|---|---|---|---|---|---|---|---|
| **Asset class risk and return characteristics** | | | | | | | | | |
| Investment grade bonds | FI | 5.6 | 4.3 | 10 | 90 | 0.40 | 3.60 | 1.26 | 0.40 |
| Municipal bonds | MFI | 7.1 | 3.6 | 10 | 90 | 0.36 | 3.24 | 0.95 | 0.24 |
| U.S. Large Company Stocks | LC | 19.2 | 8.2 | 75 | 25 | 6.08 | 2.03 | 0.84 | 0.09 |
| U.S. Small Company Stocks | SC | 28.5 | 10.0 | 85 | 15 | 8.25 | 1.46 | 0.71 | 0.53 |
| Developed market stocks | DM | 23.5 | 8.3 | 75 | 25 | 6.15 | 2.05 | 0.79 | 0.45 |
| Emerging market stocks | EM | 33.0 | 10.9 | 85 | 15 | 9.35 | 1.65 | 0.67 | 0.79 |
| Real estate | RE | 22.9 | 8.4 | 45 | 55 | 3.42 | 4.18 | 1.26 | 0.42 |
| Commodities | C | 20.0 | 4.0 | 100 | 0 | 4.30 | 0 | — | — |

| | | Income | | Capital gains[9] | | | |
|---|---|---|---|---|---|---|---|
| Asset class | | Income type | Income tax rate % $A/C^{\dagger}$ | Short-term turnover ($TO_S$) | Short-term gain tax rate % ($T_S$) A/C† | Long-term turnover ($TO_L$) | Long-term gain tax rate % ($T_L$) A/C† |
| **Taxation characteristics**[8] | | | | | | | |
| Investment grade bonds | FI | Taxable | 33/28* | 0% | 33/28 | 0% | 15/15 |
| Municipal bonds | MFI | Tax-exempt | 0/0 | 0% | 0/0 | 0% | 15/15 |
| U.S. Large Company Stocks | LC | Qualified | 15/15 | 0% | 33/28 | 0% | 15/15 |
| U.S. Small Company Stocks | SC | Qualified | 15/15 | 0% | 33/28 | 0% | 15/15 |
| Developed market stocks | DM | Qualified | 15/15 | 0% | 33/28 | 0% | 15/15 |
| Emerging market stocks | EM | Qualified | 15/15 | 0% | 33/28 | 0% | 15/15 |
| Real estate | RE | Taxable | 33/28 | 0% | 33/28 | 0% | 15/15 |
| Commodities | C | — | — | 100% | 22.2/20.2[10] | 0% | 15/15 |

†A/C = Rate in Accumulation Period/Rate in Consumption Period. *The marginal tax rate is equal to the ordinary income tax rate.

**Table 2** (*continued*)

| | | | |
|---|---|---|---|
| **Investor details** | | | |
| Years in accumulation | 30 | % of total assets in taxable accounts | 34 |
| Years in consumption | 30 | % of total assets in tax-deferred accounts | 33 |
| Intertemporal substitution rate % | 5.3 | % of total assets in tax-exempt accounts | 33 |
| Forward consumption rate % | 3.0 | Forward consumption dampening rate % | 2.75 |

Multiple 25,000 iteration simulation runs confirmed a significant improvement in the stability of present value estimates derived from the transformation of lnPV statistics across simulation runs. Table 1 presents the results of five individual simulation runs for emerging market stocks held in a tax-exempt account. This is the highest returning and most volatile asset class investment in our analysis. The variability across simulation runs in the estimates of the arithmetic means and standard deviations calculated from the raw simulated present values show the issues experienced with sampling error. Arithmetic means and standard deviations for the lnPVs across simulation runs are shown to be quite stable, as are estimates derived from transforming lnPV statistics. The distributions of the lnPVs are also shown to exhibit minimal skewness and excess kurtosis.

Having addressed the issue of estimation error, simulations can be run to produce the necessary TCPA inputs. In order to allow for simplified scaling later in the TCPA process simulations are conducted assuming a one dollar initial investment. Table 2 presents a summary of the tax-exempted asset class expectations, asset class characteristics, taxation characteristics, and investor details used to develop the inputs for our optimization.

## 4 The tax-cognizant present value frontier

Conducting a mean–variance optimization with TCPA inputs is similar to conducting a standard tax-exempted optimization except that it should consider all of the account types included in the analysis. Recall that the same asset class investment held in each of the three accounts in our analysis is treated as three distinct investments. This requires applying account-type constraints to the optimization process. Specifically, an optimizer should be appropriately constrained so that these distinct "account-type" investments are available for use only to the extent that the actual proportion of an investor's assets is located in each account type. That is, if an investor has 20 percent of his/her assets in tax-deferred accounts the optimizer should be limited to allocating no more than 20 percent to "tax-deferred" investments. As with standard optimization, additional constraints may be applied. However, for the TCPA output to be conceptually valid constraints should be consistently applied across multiple account types. The distribution of our investor's assets across account types is included in Investor details section of Table 2.

The TCPA efficient frontier we will present includes account-type constraints only. For comparison purposes the TCPA efficient frontier is presented relative to a traditional mean–variance optimized (MVO) frontier generated by applying the tax-exempted inputs used in the derivation of TCPA inputs. The optimization process employed for both TCPA and MVO portfolios creates frontiers of 100 portfolios equally spaced by standard deviation.[11] The translation of MVO portfolios to present value space is accomplished by using the asset class weights of each of the 100 frontier portfolios determined by the MVO process and then using our TCPA inputs to back into the corresponding expected PVs and expected PV standard deviations.[12]

The MVO frontier assumes the same distribution of assets across account types and that an investor implements each account identically.[13] This is not unlike how many practitioners approach implementing asset allocated portfolios for their clients. An example of how this occurs is as follows. An advisor helps an investor determine their risk tolerance and the investor, using this information, selects a "Moderate" asset allocation. The advisor then implements all of the client's accounts (i.e., taxable, Traditional IRA, and Roth IRA) with the identical "Moderate" allocation. This assumption is intended to provide a fair and objective baseline for comparison with a commonly used implementation methodology. Figure 3 presents the TCPA and MVO efficient frontiers in terms of expected PV means and expected PV standard deviations. The TCPA portfolio in this illustration was created using inputs derived from a simulation run with 25,000 iterations.

As might be expected, TCPA dominates MVO in terms of after-tax present value outcomes across most of the frontier. While the range of PV risk demonstrated by each frontier is virtually identical, there are notable differences in the expected PV means provided by the TCPA frontier. In the lower risk portion of the frontiers, TCPA PV means are notably higher than those on the MVO frontier. As PV risk increases the TCPA and MVO frontiers converge. A review and comparison of the composition of each of the frontiers will provide insight regarding this convergence and other important aspects of tax-cognizant efficiency.

Figure 4 shows the composition of the two frontiers in terms of asset classes and the accounts in which those asset class investments are held.[14] The horizontal axis of composition charts identifies specific portfolios on each frontier. TCPA frontiers are made up of 100 efficient portfolios equally spaced by present value standard deviation. MVO portfolios were optimized using tax-exempted means and standard deviations and are made up of 100 efficient portfolios spaced evenly in terms of tax-exempted standard deviations. Because of the differences in the inputs used to optimize, identically numbered portfolios from each frontier do not present equivalent risk in either tax-exempted space or present value space.

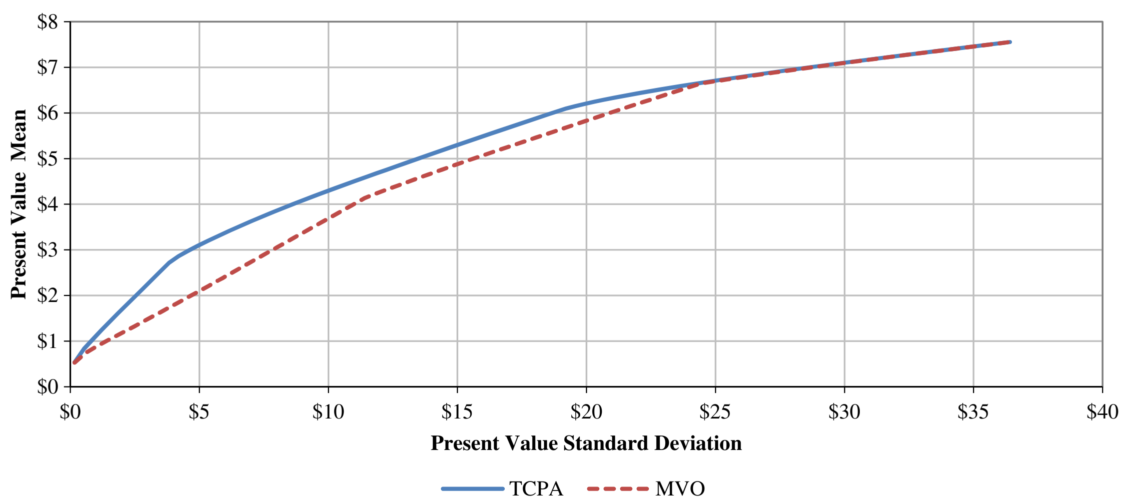


**Figure 3** A comparison of TCPA and MVO present value frontiers.

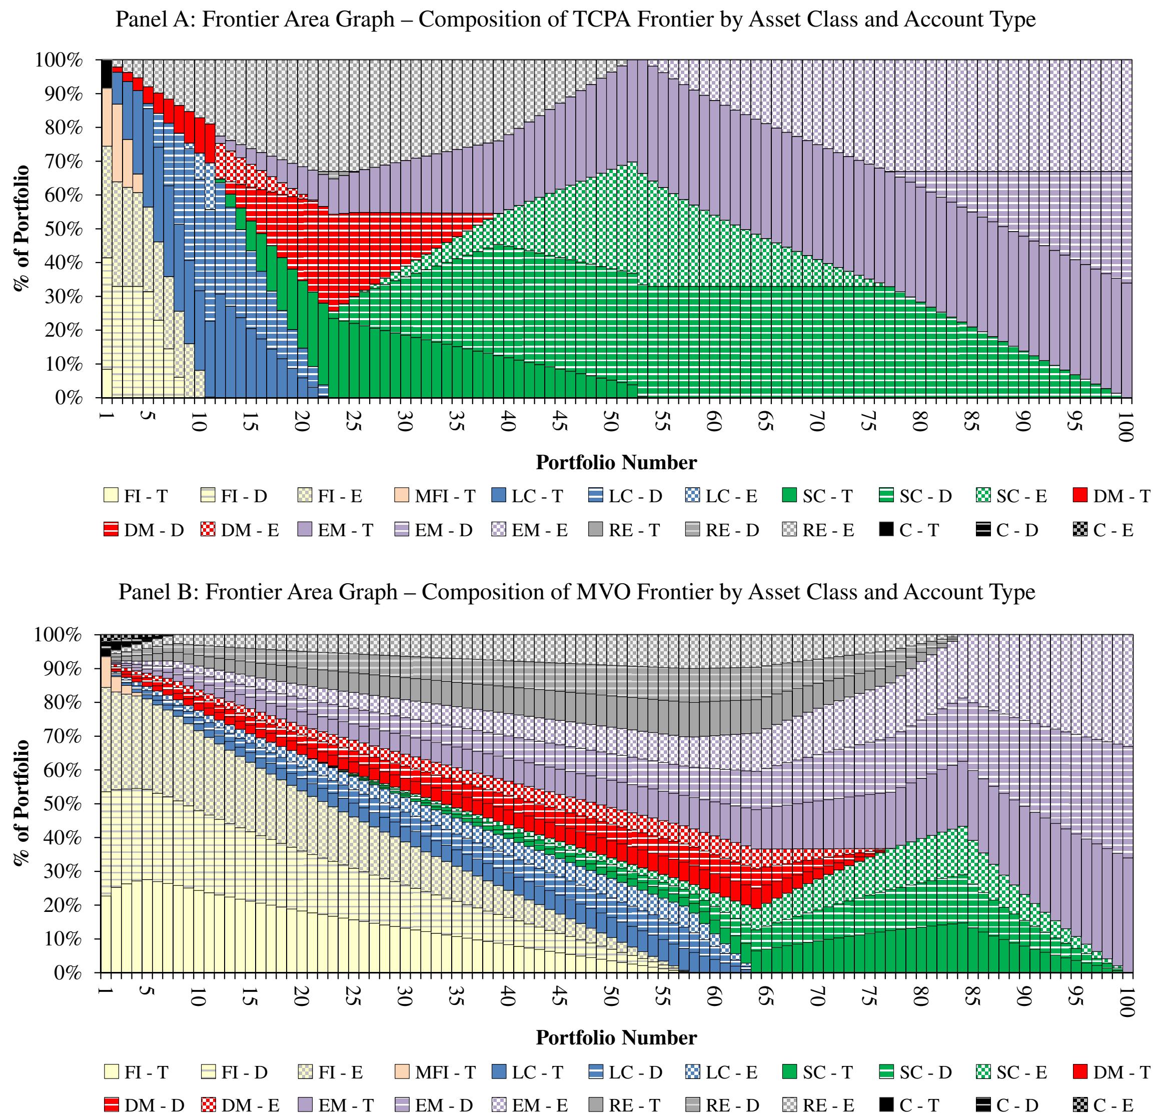


**Figure 4** The composition of TCPA and MVO efficient frontiers by asset class and account type. **FI** = Investment Grade Bonds, **MFI** = Municipal Bonds, **LC** = U.S. Large Company Stocks, **SC** = U.S. Small Company Stocks, **DM** = Developed Market Stocks ex-U.S., **EM** = Emerging Market Stocks, **RE** = Real Estate, and **C** = Commodities. Asset classes held in taxable accounts are identified with a **T** and are shown as solid colors. Asset classes held in tax-deferred accounts are identified with a **D** and are shown as colors with a horizontal pattern. Asset classes held in tax-exempt accounts are identified with an **E** and are shown as colors with a cross-hatched pattern.

An initial review of the composition of the frontiers provides the reason for the convergence of the two frontiers as risk increases. The convergence occurs at TCPA portfolio 67 and MVO portfolio 84. A comparison of these two portfolios and the portfolios that follow shows that they are composed of very similar mixes of small-cap stocks and emerging market stocks. Marginal differences in the allocations and locations of these two asset classes result in small differences in expected present values. The composition of both frontiers at their highest point is identical as they are entirely made up of emerging market stocks, the asset class with the highest return, distributed according to the percentage of our investor's assets that can be located in each of the three account types.

Other aspects of the optimizer's preferences for investments are revealed through a further review of the composition of the frontiers. One of the most notable differences between frontiers in terms of asset class preferences can be seen in the amount of investment grade bonds incorporated. As might be expected with MVO, the amount of investment grade bonds slowly decreases as portfolio risk increases. With TCPA, fixed income falls out early in the frontier. This can partially be attributed to the fact that the expected total return for fixed income (4.3 percent) is below the intertemporal rate of substitution (5.3 percent) we used to discount the consumption cash flows.

The apparent early exit of investment grade bonds from the TCPA frontier can also be attributed to the large differences between the standard deviations of present values for the safest and riskiest assets included in the analysis. The volatility of present values is a function of the risk and return characteristics of the asset classes included in the analysis and the length of the investment horizon. Consider a simulation for emerging market stocks over a 60-year investment horizon where the asset class rarely experienced a major drawdown. The compounded growth of a dollar and the consumption cash flows provided in such a simulation might be phenomenal over such a long investment time-horizon, especially in a tax-exempt account. However, because emerging markets is also the riskiest asset class in the analysis, there are simulation outcomes where the asset class experienced multiple significant drawdowns. The large variability in these possible outcomes indicates that the standard deviation of the present values for emerging market stocks can be expected to be large.

In the simulation run used to produce our TCPA inputs investment grade bonds held in a taxable account had a PV standard deviation of \$0.13, while emerging market stocks held in a tax-exempt account had a PV standard deviation of \$45.00. Recall that the frontier is composed of portfolios that are equally spaced according to PV standard deviation. This means that portfolio spacing is largely determined by the riskiest and highest returning asset class in the analysis and the length of the investment horizon. In this instance, the standard deviation interval between portfolios is approximately \$0.37. As such, less-risky assets like fixed income appear to fall out quickly from portfolios in the presented frontier. The standard deviation intervals between portfolios in a TCPA frontier decrease as the investment time horizon shortens.

With regard to optimal asset location, a few general observations can be made. The TCPA process prefers to place asset classes which produce a large portion of their returns from income that is taxed at the investor's marginal tax rate in tax-deferred or tax-exempt accounts. Real estate is a good example of this preference as it is never presented in the TCPA frontier in a taxable account. Another observation regarding TCPA asset location preferences is that as an investor progresses through the frontier from left to right, assets generally enter portfolio compositions in taxable accounts then move their way into tax-deferred and finally into tax-exempt accounts. This is because the tax-impact experienced by different accounts types affects not only the expected present values, but also the standard deviation of the present values. The more tax-efficient the account the higher the present values and PV standard deviations. The optimizer first makes the best use of lower risk investments and then shifts into riskier investments as it moves up the frontier. This has implications for asset location approaches currently being practiced by financial advisors who use generalized rules for asset location based primarily on achieving the highest after-tax return provided by asset classes within different accounts. Optimal asset location is a function of both risk and after-tax return.

### 4.1 Frontier resampling

Before final portfolio recommendations can be made using TCPA, it is necessary to address

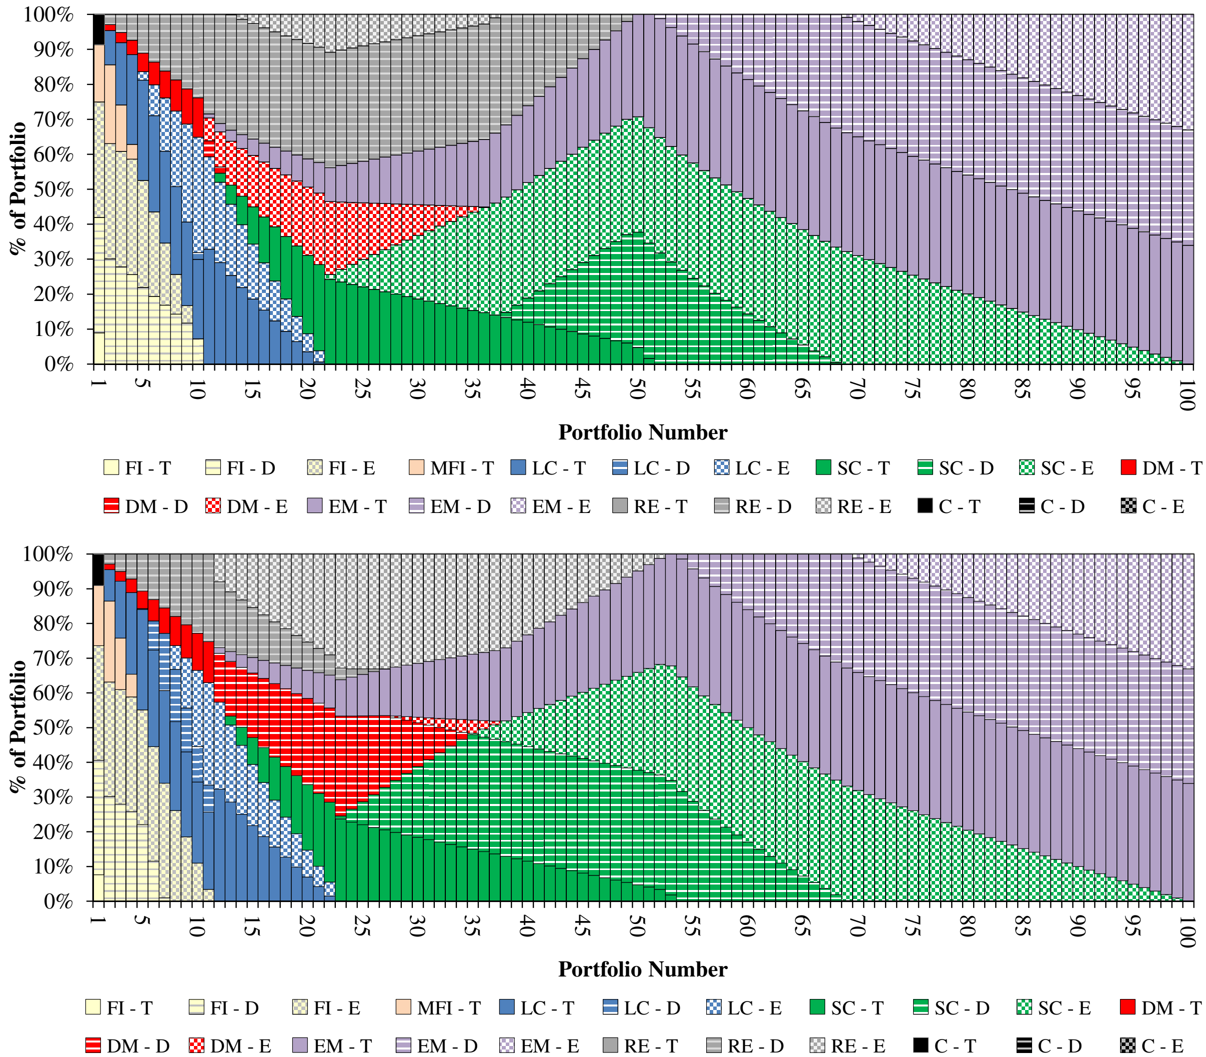


**Figure 5** Comparison of the composition of two tax-cognizant present value frontiers produced using the same tax-exempted inputs. **FI** = Investment Grade Bonds, **MFI** = Municipal Bonds, **LC** = U.S. Large Company Stocks, **SC** = U.S. Small Company Stocks, **DM** = Developed Market Stocks ex-U.S., **EM** = Emerging Market Stocks, **RE** = Real Estate, and **C** = Commodities. Asset classes held in taxable accounts are identified with a **T** and are shown as solid colors. Asset classes held in tax-deferred accounts are identified with a **D** and are shown as colors with a horizontal pattern. Asset classes held in tax-exempt accounts are identified with an **E** and are shown as colors with a cross-hatched pattern.

the issue of instability in the composition of tax-cognizant frontiers. The interaction between random return sequences, taxation dynamics, and the large number of variables that must be specified, produces frontiers that, while virtually identical in terms of PV risk and returns outcomes, demonstrate some variability in the composition of frontiers. Frontiers generally demonstrate only slight differences in terms of the asset class allocations. However, there are more notable differences in the location of those asset classes amongst account types across simulations. Figure 5 shows two TCPA frontiers developed using identical tax-exempted inputs. Differences in the locations of large company stocks, real estate, developed market stocks, and small company stocks are apparent. The differences in frontiers occur because simulation runs sometimes result in present value outcomes that make asset class investments very close substitutes. Variability across frontier outcomes persisted even when inputs were derived from simulations that included 200,000 iterations. This indicates that the production of stable frontier compositions requires the use of a frontier resampling process where several frontiers are produced and outcomes from multiple frontiers are averaged. The resampling process used depends on the portfolio selection approach chosen. More details regarding resampling and portfolio selection approaches are provided later in this paper.

### *4.2 Portfolio selection*

We have demonstrated that the TCPA methodology can provide investors with the intended benefit of maximizing expected after-tax wealth outcomes through both asset allocation and asset location. However, an impractical aspect of the tax-cognizant frontier we have presented, that may not be evident initially, is found in the difficulty an investor might have in selecting an appropriate portfolio, given the fact that it is presented in terms of PV mean and PV standard deviation. While it is evident that an investor would want a higher present value than a lower present value, it is unclear how that directly translates into an investor's ability to plan for future wealth consumption. Furthermore, the mean of the lognormally distributed present values is a poor estimator of expected outcomes. Even less intuitive is the notion of the standard deviation of present values. With lognormally distributed present value outcomes, the act of conceptualizing risk using standard deviation is a difficult, if not nebulous, proposition. In short, portfolio selection using expected present value outcomes in this manner is not intuitive. For TCPA to be useful to investors, it will be necessary to provide a more intuitive framework that allows an investor to easily conceptualize and form preferences amongst different portfolios on a tax-cognizant efficient frontier.

## 5 Tax-cognizant cash flow–confidence level frontiers

Portfolio selection for an investor should reflect the objective of seeking to maximize expected after-tax wealth where wealth is defined as the cumulative value of the after-tax cash flows that an investment can provide over the investor's lifetime. The important aspects of this objective for portfolio selection are determining the amount of the after-tax cash flows an investor can expect and the likelihood of achieving those cash flows.

### *5.1 Transforming the present value efficient frontier*

Before transforming our present value outcomes to more useful parameters for portfolio selection we should first review how those present values were determined. Recall that the process of deriving the inputs for our TCPA centered on determining the after-tax consumption cash flows provided by asset class investments. These cash flows were then discounted at an interest rate that described our investor's intertemporal consumption preferences. What makes the present value results we have presented unintuitive is that they represent the sum of a series of annual cash flows. A more intuitive presentation of tax-cognizant outcomes would be to present them in terms of inflation adjusted (real) average annual cash flows.

While we do not know the specific amounts of each of the annual consumption cash flows over the consumption period, we do know the present values of the cash flows, the discount rate used, the length of the accumulation period, and the number of annual cash flows in our consumption period. If we grow the present value of cash flows at the rate of intertemporal substitution (our discount rate) over the accumulation period we are left with the sum of discounted cash flows at the beginning of the consumption period. With this information we can determine the average annual after-tax cash flow.[15] Because the intertemporal rate of substitution used to discount cash flows incorporates an inflation component, the average annual after-tax cash flow we calculate will already be expressed in real terms.

The investor can now be presented with the average annual real after-tax cash flow they can expect from each TCPA frontier portfolio. Because the TCPA inputs used were derived from simulations that assumed an initial investment of one dollar, we can scale outcomes to match our investor's circumstances by multiplying expected present values by the current value of our investor's assets.

To address the unintuitive nature of the standard deviation of lognormally distributed outcomes and provide the investor with a measure that indicates the likelihood of receiving the average annual real after-tax cash flow, we again make use of the relationship between present values and their normally distributed logs.

In a normal distribution, the mean is the central point of the distribution where 50 percent of outcomes are expected to fall to the left of the mean and 50 percent are expected to fall to the right of the mean. Because of this symmetry, the mean is equal to the median (the middle value of observations). In a lognormal distribution, present value outcomes are not symmetrically distributed around the mean, they are skewed to the right. In this case, the mean and median are different and the median is the appropriate measure of central tendency. This is because the mean is pulled higher by present value outcomes in the skewed right tail of the lognormal distribution while the median remains much less affected by the distribution's skew and outliers. Figure 6 shows just how far the means of frontier portfolios with lognormally distributed present value outcomes are pulled from their median values. The median, a robust estimator of central tendency, is a much more useful measure for portfolio selection when outcomes are expected to be lognormally distributed.

The median of a portfolio's lognormally distributed present values is determined by transforming the expected, or mean, present value of a portfolio to the expected value of the lnPVs, which are normally distributed. Recall that as part of deriving TCPA inputs we transformed the mean of the simulated lnPVs to the mean of present values. Here we do the reverse. We transform the mean of present values to mean of lnPVs. Because the expected value, or mean, is the same as the median for normal distributions, we transform the lnPV median back to a present value by taking its exponential ($e^{\ln PV}$) to identify the median present value.

We can use this same method to identify other points on the normally distributed lnPV

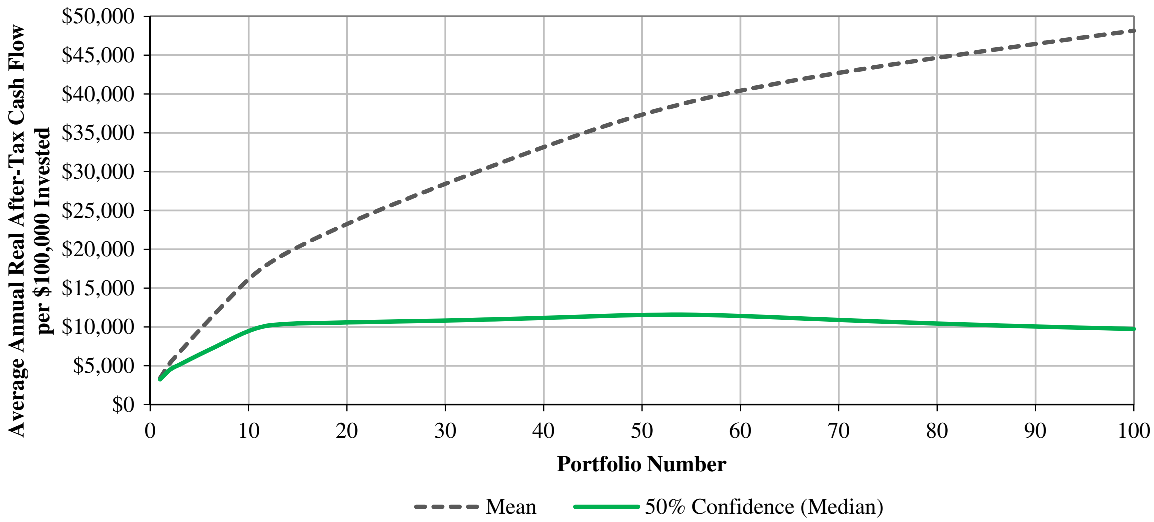


**Figure 6** A comparison of the mean and median average annual real after-tax cash flow outcomes of frontier portfolios.

distribution and transform them back to lognormal present values. Of particular interest in determining a more intuitive method of selecting portfolios from a tax-cognizant efficient frontier are left-tail confidence levels that can be identified for a normal distribution.

An investor concerned with the Value-at-Risk (VaR) of a portfolio identifies a specific left-tail confidence level that will be used for calculating that value. We are also concerned with left-tail events as they are the most critical to an investor's future consumption of wealth. Using the means and standard deviations of our lnPVs we can identify values at specific confidence levels[16,17] for the normally distributed lnPVs and then transform them back to present values by taking the exponential of the lnPV ($e^{\ln PV}$). The present values are then converted to average annual real after-tax cash flows. Similar to the information provided by a portfolio's VaR, identifying the expected present values at specific confidence levels provides the investor with a minimum average annual real after-tax cash flow outcome that might be expected from a portfolio, given a particular level of confidence.

The use of confidence levels for portfolio selection allows for a variety of transformations of the PV mean–PV standard deviation tax-cognizant efficient frontier. Here we present two transformations that result in two distinct approaches to portfolio selection.

The first is the most direct transformation where cash flow outcomes for PV mean–PV standard deviation frontier portfolios are shown at the median, or 50 percent confidence level, along with the cash flow outcomes at other confidence levels. The $X$-axis of the frontier identifies frontier portfolio numbers. The $Y$-axis identifies the average annual real after-tax cash flow. We will call this transformation the cash flow–confidence level approach to portfolio selection and it requires the presentation of multiple cash flow–confidence level frontiers. Showing median cash flow outcomes alone does not provide an understanding of the risk an investor faces for seeking a higher median value. Frontiers for other confidence levels are necessary to provide a reference for the risk–return trade-off. For illustrative purposes we have selected two additional confidence levels: 75 percent and 95 percent.

A cash flow–confidence level frontier is produced using the following resampling process. Multiple tax-cognizant present value frontier samples are produced. The compositions of corresponding individual frontier portfolios from each of the sample frontiers are averaged across frontiers. For example, the composition of resampled portfolio number five is the average of the compositions of the number five portfolios from all the sample frontiers produced. Portfolio median present values and confidence levels are then calculated by using large samples of constituent present value outcomes produced by aggregating present values from all of the simulations used to produce all of the sample frontiers. That is, if a resampled frontier is produced using 250 individual sample frontiers, each of which required 1,000 simulation runs to produce, then average present values, standard deviations, and correlations for portfolio investments are calculated based on samples of 250,000 present values per specific investment. This information is then used to determine median present values and present values for given confidence levels for all resampled frontier portfolios.[18]

Panel A in Figure 7 shows tax-cognizant cash flow–confidence level frontiers at 50 percent confidence (the median), 75 percent confidence, and 95 percent confidence. These resampled frontiers were created by averaging 250 present value frontiers developed using simulations with 1,000 iterations each. Presenting tax-cognizant outcomes at different confidence levels allows an investor to understand the implications of seeking a higher median outcome. In this instance portfolio one provides the maximum cash flow outcome at the highest confidence level presented (95 percent). However, an investor selecting this portfolio must be willing to give up significant cash flow at lower confidence levels. Less risk-averse investors might be compelled to select portfolio 12 where the 75 percent confidence level reaches its maximum. This presents attractive outcomes at both 95 and 50 percent confidence levels. An even less risk-averse investor might select portfolio 53 where the 50 percent confidence level reaches its maximum. In doing so, the investor must be willing to give up approximately one-sixth of the maximum cash flow at the 75 percent level and over half of the maximum cash flow at the 95 percent level. No rational investor would select portfolios above portfolio 53 as cash flow outcomes decrease at all confidence levels from that point forward presenting less attractive trade-offs than other portfolios on the frontier. Panel B in Figure 7 shows the composition of the resampled tax-cognizant frontier.

The second transformation of the PV mean–PV standard deviation tax-cognizant efficient frontier focuses on identifying frontier portfolios that maximize specific confidence levels. This requires the determination of average annual real after-tax cash flow outcomes for all 100 present value frontier portfolios at all confidence levels between 95 percent and 50 percent in one percent increments. Portfolios that maximize the cash flow outcomes for each confidence level are selected. These cash flow maximizing portfolios form a new tax-cognizant frontier. The $X$-axis of this frontier identifies confidence levels from 95 percent to 50 percent. The $Y$-axis identifies the average annual real after-tax cash flow. We will call this transformation the maximum cash flow–confidence level approach to portfolio selection. As with the previous approach, additional cash flow–confidence level frontiers are provided as a reference for the risk–return trade-off. For this approach we provide the upper and lower confidence levels: 50 percent (the median) and 95 percent.

A maximum cash flow–confidence level frontier is produced using the following resampling process. Multiple maximum cash flow–confidence

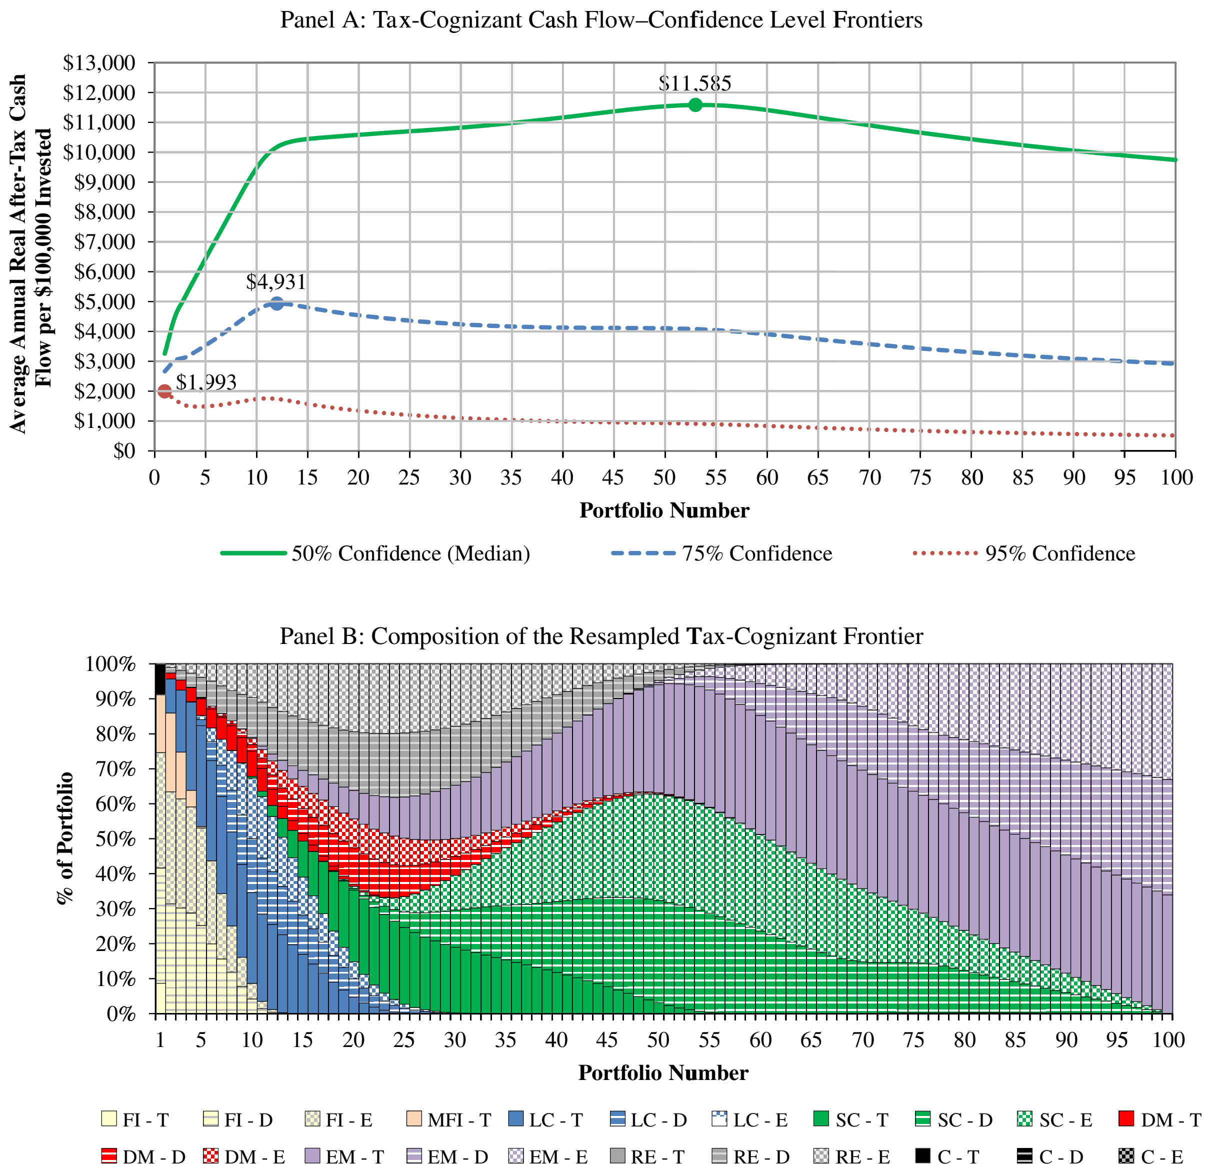


**Figure 7** Cash flow–confidence level frontiers for an investor 30 years from retirement planning for 30 years in retirement. Panel A shows multiple cash flow–confidence level frontiers. Maximum cash flow outcomes for each confidence level are also identified. Panel B shows the composition by asset class and account type of the resampled tax-cognizant frontier that corresponds to the confidence level frontiers. **FI** = Investment Grade Bonds, **MFI** = Municipal Bonds, **LC** = U.S. Large Company Stocks, **SC** = U.S. Small Company Stocks, **DM** = Developed Market Stocks ex-U.S., **EM** = Emerging Market Stocks, **RE** = Real Estate, and **C** = Commodities. Asset classes held in taxable accounts are identified with a **T** and are shown as solid colors. Asset classes held in tax-deferred accounts are identified with a **D** and are shown as colors with a horizontal pattern. Asset classes held in tax-exempt accounts are identified with an **E** and are shown as colors with a cross-hatched pattern.

level frontier samples are produced. The compositions of corresponding confidence level maximizing frontier portfolios from each of the sample frontiers are averaged across frontiers. For example, the composition of the resampled portfolio that maximizes the 80 percent confidence level is the average of the compositions of the portfolios that maximize cash flows at the 80 percent confidence level in each of the sample frontiers produced. As with the previous resampling process, portfolio median present values and confidence levels are then calculated by using large samples of constituent present value outcomes produced by aggregating present values from all of the simulations used to produce all of the sample frontiers. Therefore, median and specific confidence level outcomes for frontier portfolios are calculated using average present values, standard deviations, and correlations derived from a significantly larger sample than those used to create any single sample frontier. With long investment horizons, such as the one used in our example, this procedure can sometimes result in a situation where the estimated present value outcome for a resampled portfolio that maximizes a specified confidence level is lower than that provided by the adjacent portfolio with a higher confidence level. This generally occurs in segments of the frontier where there is a quick transition from lower risk assets to higher risk assets. In these instances standard deviation increases at a much higher rate than expected present values. This results in lower values for left-tail confidence levels. These situations cease to occur as the investment horizon is shortened.

Panel A in Figure 8 shows the maximum cash flow–confidence level frontier. This resampled frontier was created by averaging 250 maximum cash flow–confidence level frontiers developed using simulations with 1,000 iterations each. Presenting portfolios that maximize cash flow outcomes for a wide range of confidence levels allows an investor to select portfolios based on the trade-off between average annual real after-tax cash flow and the confidence level in achieving those outcomes. An investor seeking greater confidence in cash flow outcomes must be willing to give up cash flow. Conversely, an investor needing greater cash flow must be willing to accept less confidence in achieving those outcomes. Including information regarding cash flow outcomes at the 50 percent and 95 percent confidence levels allows the investor to understand the implications of seeking greater confidence in cash flows. At the 95 percent confidence level an investor must be willing to accept a median cash flow outcome that is approximately two-thirds lower than what might result from selecting the portfolio at the 92 percent confidence level. A less risk-averse investor selecting the portfolio at the 50 percent confidence level has a greater median cash flow outcome but must accept a one-in-twenty chance of receiving approximately half of the cash flow provided by the 75 percent confidence level portfolio at the 95 percent confidence level.

An obvious aspect of the maximum cash flow–confidence level frontier that can be seen in Panel B of Figure 8 and one that should be addressed is the fact that the frontier is composed mostly of equity asset classes. This is a function of the length of the investment horizon in our analysis. Given a 60-year investment horizon (30 years in accumulation and 30 years in consumption) the portfolios identified by the process maximize expected cash flows even when considering the higher volatility that might be experienced with the higher weighting to equities. The frontier was developed based on several thousand simulated expected present value outcomes. Some of these present values were made up of lifetimes that experienced large drawdowns and provided less favorable outcomes while other lifetimes did not see large drawdowns and provided much more favorable outcomes. This is one of the benefits

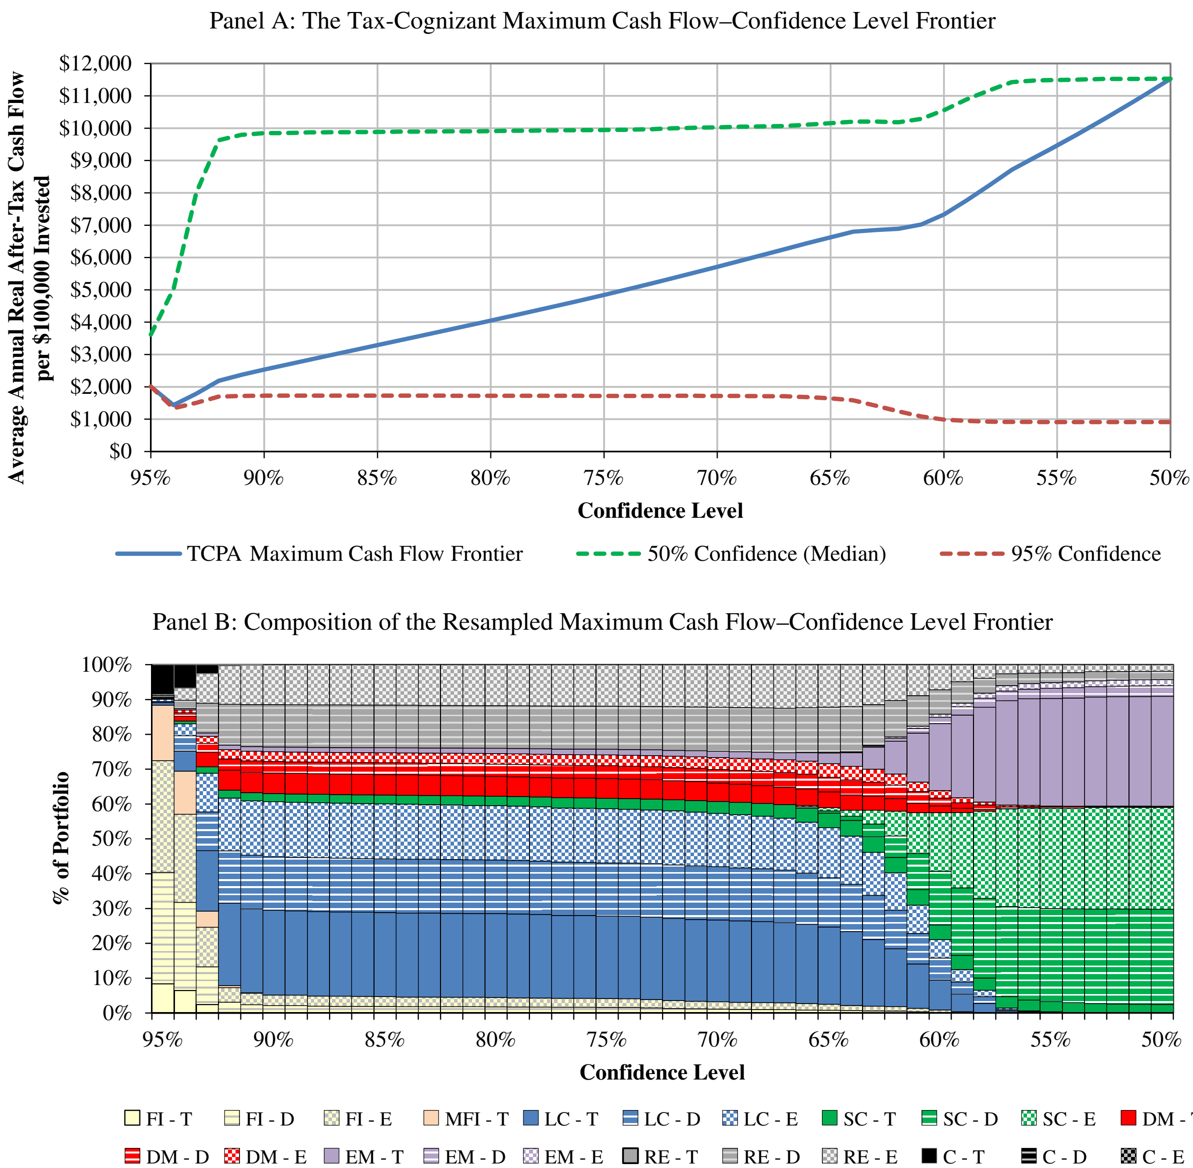


**Figure 8** The maximum cash flow-confidence level frontier for an investor 30 years from retirement planning for 30 years in retirement. Panel A shows the maximum cash flow–confidence level frontier with cash flow outcomes of frontier portfolios at the 50 percent and 95 percent confidence levels. Panel B shows the composition of the maximum cash flow-confidence level frontier by asset class and account type. **FI** = Investment Grade Bonds, **MFI** = Municipal Bonds, **LC** = U.S. Large Company Stocks, **SC** = U.S. Small Company Stocks, **DM** = Developed Market Stocks ex-U.S., **EM** = Emerging Market Stocks, **RE** = Real Estate, and **C** = Commodities. Asset classes held in taxable accounts are identified with a **T** and are shown as solid colors. Asset classes held in tax-deferred accounts are identified with a **D** and are shown as colors with a horizontal pattern. Asset classes held in tax-exempt accounts are identified with an **E** and are shown as colors with a cross-hatched pattern.

of using confidence levels for portfolio selection. At the 95 percent confidence level we would expect 95 percent of the cash flow outcomes to be the cash flow indicated or better. In fact, the cash flow indicated at the 95 percent confidence level would be amongst the worst to be expected. As the time horizon shortens, maximum cash flow–confidence level frontier efficient frontiers incorporate increasing amounts of less volatile assets like investment grade fixed income, especially as part of the higher confidence segments of frontiers.

An important ancillary benefit of presenting frontiers using confidence levels can be seen through a comparison of the composition of frontier portfolios and the expected after-tax cash flow outcomes they provide. A review of the composition of the portfolios for the 90 percent through 75 percent confidence levels shows them all to be very similar. However, expected after-tax cash flows provided by the portfolios increase from approximately $2,534 per year to $4,842 as confidence decreases. This highlights that the maximum cash flow–confidence level portfolio selection approach is useful for providing guidance on both portfolio selection and also on the amount an investor might choose to consume, given an investor's need for confidence in cash flow outcomes.

### 5.2 *The benefits of tax-cognizant portfolio analysis*

Figure 9 presents a comparison of the average maximum cash flow–confidence level outcomes from 250 TCPA and MVO sample frontiers produced using 1,000 iterations each. This demonstrates that the TCPA process is able to identify portfolios that provide investors with real after-tax annual cash flow outcomes that are notably more desirable than those provided by an MVO process using tax-exempted inputs. In this instance, improvements to cash flow outcomes range from 3 percent to 54 percent across the maximum cash flow–confidence level frontier. A careful review of the differences in outcomes also reveals that the greatest benefits generally occur at the higher confidence/lower risk segments of the

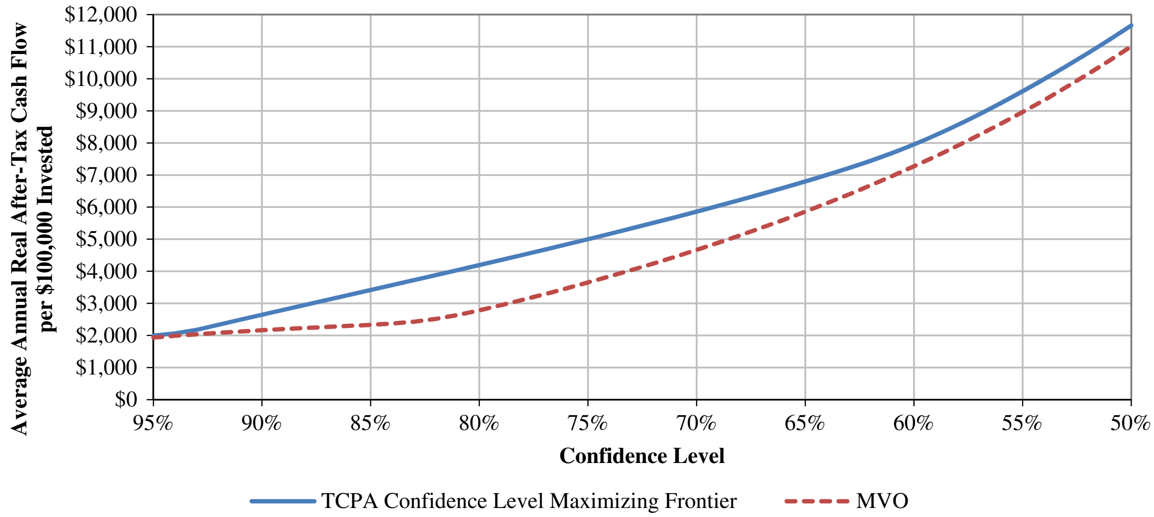


**Figure 9** A comparison of the average maximum cash flow–confidence level outcomes from 250 TCPA and MVO sample frontiers produced using 1,000 iterations each.

frontier. These improvements are accomplished through both asset allocation and asset location.

It should be understood that the expected benefits achieved by TCPA are a function of the parameters used in the analysis. It is intuitive that higher tax rates will result in greater expected benefits. Longer investment time horizons also result in greater benefits due to the effects of tax-deferred compounding. The tax-exempted inputs used also affect the degree to which TCPA is expected to benefit investors. Changes in expected returns, standard deviations, and correlations can all affect the benefits expected from TCPA. The results presented in this paper are modest relative to other completed analyses. Our research identified instances where unconstrained TCPA frontiers demonstrated improvements over MVO frontiers of over 140 percent and constrained TCPA frontiers showed improvements of over 90 percent.

### *5.3 Risk considerations for tax-cognizant portfolio analysis*

The objective of TCPA is to maximize after-tax wealth outcomes where wealth is represented by a series of after-tax cash flows. These cash flows occur over an investment horizon that generally includes more than one periodic investment return outcome. In this context, the risk defined by the TCPA process is that of the expected cumulative effect of a sequence of investment return outcomes on a sequence of after-tax cash flows. It describes the characteristics of cash flow outcomes over complete horizons rather than at any given point within those horizons. An intellectual understanding of this fact may be of little consolation for investors who are forced to contend with the intra-horizon risks posed to consumption cash flows by periodic portfolio outcomes as they progress through those horizons. Therefore, successful portfolio selection from a TCPA frontier should consider the risks that occur over the horizon as well as risks that might be expected to occur within the horizon. While our presentation of TCPA has focused on risk in terms of after-tax cash flows, the use of the traditional portfolio return standard deviation should be an integral part of the determination of an appropriate portfolio. Ideally, a well selected TCPA portfolio will balance confidence in acceptable after-tax cash flow outcomes with tolerable levels of intra-horizon portfolio return volatility.

## 6 TCPA, time-dependence, and tax-cognizant investing

The research conducted in the development of the TCPA process revealed a very important difference between TCPA and a traditional mean–variance portfolio analysis. The results of a tax-exempted mean–variance portfolio analysis are time-independent while the results of a TCPA are time-dependent. In other words, mean–variance produces the same results regardless of the time horizon used, while the results of a TCPA differ with changes in the time horizon.

To illustrate how changes in the investment time horizon affect TCPA results we present a series of portfolios produced by incrementally conducting TCPAs over our investor's lifetime. All of the specifications used to create the TCPA frontier presented earlier will remain the same except that portfolios will be drawn from resampled frontiers derived from 500 sample frontiers. The lengths of the accumulation and consumption periods are adjusted in accordance with a progression through an investor's lifecycle. We assume that the investor starts with the same expectation of 30 years in accumulation and 30 years in consumption. We then advance through the investor's lifecycle in one year increments. Therefore, the years expected in retirement will decrease only after the accumulation period has

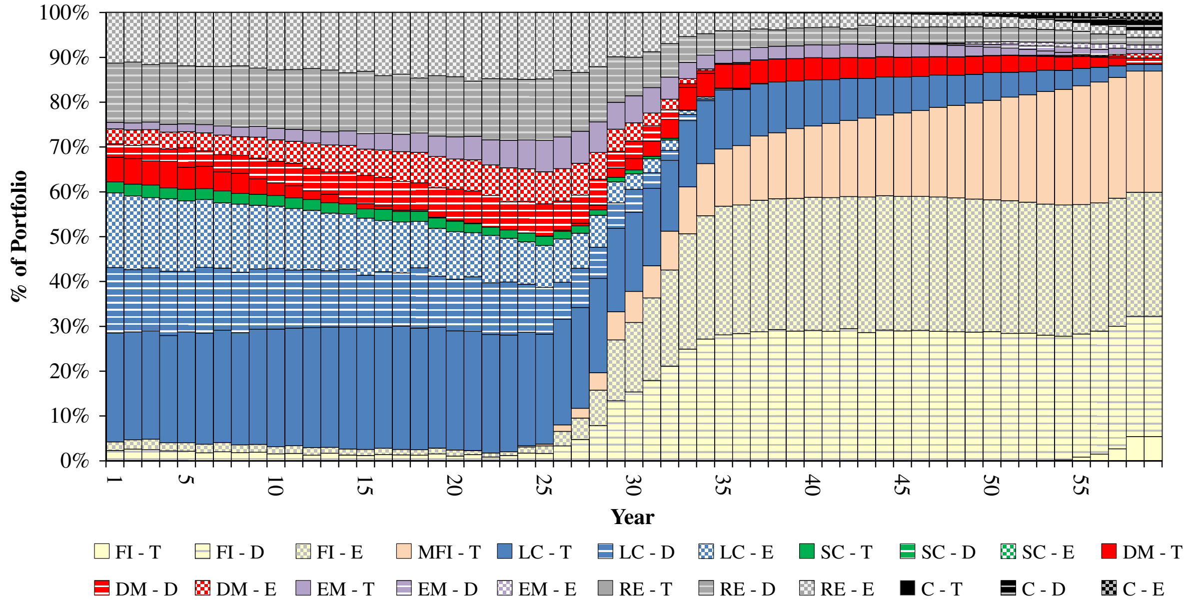


**Figure 10** An example of a series of 82 percent confidence level TCPA portfolios produced using different investment horizons. The portfolio at the left-most part of the panel assumes that an investor expects to be in accumulation for 30 years and retirement for 30 years. Progressing from left-to-right the years in accumulation decrease until year 30 where years in consumption then begin to decrease until the expected end of retirement. The portfolios were developed using the TCPA analysis assumptions presented in Table 2. **FI** = Investment Grade Bonds, **MFI** = Municipal Bonds, **LC** = U.S. Large Company Stocks, **SC** = U.S. Small Company Stocks, **DM** = Developed Market Stocks ex-U.S, **EM** = Emerging Market Stocks, **RE** = Real Estate, and **C** = Commodities. Asset classes held in taxable accounts are identified with a **T** and are shown as solid colors. Asset classes held in tax-deferred accounts are identified with a **D** and are shown as colors with a horizontal pattern. Asset classes held in tax-exempt accounts are identified with an **E** and are shown as colors with a cross-hatched pattern.

ended. Figure 10 shows the results of this analysis and the changes to the composition over time for 82 percent maximum confidence level portfolios.

While this series of portfolios could be construed as a predetermined investing glide path, there are several dynamic variables that are part of the TCPA process that make it infeasible to use it in such a manner. The volatility in the capital markets over time will change the distribution of an investor's assets across account types. Contributions to and withdrawals from accounts can also change the distribution of assets across accounts. The amount of unrealized gains within an investor's taxable portfolio also varies with market volatility. The changes in these dynamic variables require that a new TCPA be conducted in order to optimize the investor's updated state of taxation. Including unrealized gains as part of ongoing TCPA analyses ensures that portfolios recommended by the process respect the tax implications of shifting assets in taxable accounts. Analyses may need to be appropriately constrained in order to avoid unnecessary or undesirable taxable events.

Along with changes that result from volatile capital markets, we have demonstrated that the time-dependent nature of the TCPA process also requires that investors review their allocation on a regular basis to achieve the benefits of tax-cognizant investing and to maintain desired confidence levels in cash flow outcomes over time. This would not present an issue for an investor who is engaged in the process of planning for and attempting to maximize expected retirement cash flows. We believe that the TCPA process would be conducive to constructive conversations between advisor and investor as they would be presented with an expected average annual cash flow on a regular basis. An investor presented regularly with such information is much more likely to make the necessary decisions to achieve a successful retirement outcome and, consequently, much more likely to be prepared for retirement. Portfolio selection using confidence levels and expected cash flow outcomes also allows an investor to opportunistically adjust the desired confidence level once specific expected retirement goals are achieved to a desired degree of confidence.

For investors who are less engaged in their retirement planning, the time-dependent characteristics of TCPA could present a mismatch in the risks presented by TCPA portfolios and the investor's willingness to bear those risks. For these situations we describe two approaches for using TCPA that may be effective.

The first approach to dealing with less engaged investors would be a constant confidence level strategy. For this approach an investor would select a desired level of confidence for cash flow outcomes. The TCPA portfolio would then automatically adjust over time in order to continue maximizing outcomes for the selected confidence level. An example of this is an investor who initially selects a 75 percent confidence level portfolio. Every year the portfolio allocation would be adjusted to match the 75 percent confidence level portfolio that was appropriate for the updated time horizon. This would be similar to the process of selecting higher or lower risk portfolios in use today. Because TCPA portfolios for long investment horizons tend to favor equities, TCPA practitioners would need to be sensitive to the intra-horizon portfolio risks that might be presented by such recommendations.

The second approach to dealing with less engaged investors would be an increasing confidence level approach. For this approach an investor would initially select a desired level of confidence for cash flow outcomes. The TCPA portfolio would then automatically shift to portfolios that maximize increasing confidence levels over time. The portfolio is gradually shifted towards the portfolio that maximizes the highest confidence level desired by the investor. Reaching this maximum could be programmed to occur at an investor's expected retirement date or at any time specified by the investor. An example of this is an investor who elects to begin investing with a 50 percent confidence level portfolio and desires to shift to an 80 percent confidence level portfolio as she approached retirement. In this instance the portfolio would incrementally shift to higher confidence level portfolios until the 80 percent confidence level portfolio is reached at the planned retirement date. Portfolios would then follow a constant confidence approach.

These approaches might seem similar to target-date or lifecycle portfolios used today that follow a pre-specified glide path to adjust portfolios over time. However, a distinct difference with these approaches is that future portfolios cannot be pre-specified due to the variable nature of the necessary inputs. Portfolios used in any given year require conducting a new TCPA with the dynamic variables that define the investor's state of taxation at that point in time. This could be viewed as the process of following a dynamic tax-cognizant glide path.

## 7 Conclusion

We have introduced a Tax-Cognizant Portfolio Analysis methodology intended to allow investors to make more effective decisions regarding the maximization of after-tax wealth. The application of the TCPA process was detailed beginning with tax-exempted inputs through the creation of maximum cash flow–confidence level efficient frontiers. The benefits of applying TCPA relative to traditional MVO approaches using tax-exempted inputs were shown to be notable. The analysis presented in this paper demonstrated improvements in expected after-tax cash flow outcomes ranging from 3 percent to 54 percent across the maximum cash flow–confidence level frontier with the greatest benefits being seen in the higher confidence/lower risk segments of frontiers.

We contribute to asset allocation literature in three key areas. First, we detail the TCPA methodology that effectively deals with investment and consumption-based illiquidities and incorporates the impact of taxation into the process of constructing portfolios that seek to maximize after-tax wealth for given levels of risk. This included providing detail on the use of present values for portfolio optimization and overcoming the issue of estimation error in determining the means, standard deviations, and correlations of a series of simulated present value outcomes. Second, the lognormal distribution of present value outcomes necessitates that tax-cognizant portfolios be presented using different, more intuitive, risk and return metrics for portfolio selection. We detail two transformations of present value efficient frontiers and present two distinct approaches to tax-cognizant portfolio selection. The cash flow–confidence level approach allows investors to select portfolios based on real after-tax periodic cash flow outcomes at various confidence levels. The maximum cash flow–confidence level approach allows investors to select portfolios that maximize real after-tax periodic cash flows for a given probability of achieving those cash flows. Finally, because the results of a TCPA are dependent on the length of the investment horizon we present the use of dynamic portfolio glide paths that automatically balance tax efficiency with the risk of not achieving desired after-tax wealth outcomes.

## Appendix

*The Blay–Markowitz taxation model*

The taxation model allows for the determination of after-tax cash flow outcomes from various account types. Much of the model's complexity is required to deal with taxation that is pertinent only to taxable accounts. The model assumes consumption occurs at the beginning of each period and incorporates the fractional wealth consumption model discussed in this paper. Consumption is determined by equations (A.1)–(A.3), which can be modified to accommodate alternative consumption methodologies. When using the model with tax-deferred accounts all tax rates should be set to zero except the marginal tax rate. When using the model with tax-exempt accounts all tax rates should be set to zero including the marginal tax rate.

Note: In the formulas hereafter, the indicator function notation style $\mathbb{1}(X)$ is equivalent to the expression:

$$[X] = \begin{cases} 1 & \text{if X is true;} \\ 0 & \text{otherwise.} \end{cases}$$

Definition of Variables:

| | |
|---|---|
| $n_A$ | Expected number of years remaining in accumulation period |
| $n_C$ | Expected number of years remaining in consumption period |
| $W_T$ | Total wealth |
| $W_A$ | After-tax wealth |
| $W_U$ | Untaxed wealth |
| $R_I$ | Income return |
| $R_C$ | Capital return |
| $TO_S$ | Short-term turnover ratio (*value between 1 and 0*) |
| $TO_L$ | Long-term turnover ratio (*value between 1 and 0*) |
| $T_M$ | Marginal tax rate |
| $T_I$ | Tax Rate – income |
| $T_S$ | Tax Rate – short-term capital gain |
| $T_L$ | Tax Rate – long-term capital gain |
| $T_N$ | Net taxes |
| $L$ | Loss carry-forward |
| $L_U$ | Loss carry-forward used to offset consumption |
| $C_A$ | Consumption from after-tax wealth |
| $C_U$ | Consumption from untaxed wealth |
| $F$ | Forward consumption rate |
| $D$ | Forward consumption dampening rate |
| $cmf$ | Consumption modification factor |
| $AC$ | After-tax consumption |
| $\mathbb{1}(n_A = 0)$ | Tests whether the investor is currently in accumulation or consumption |
| $\mathbb{1}(R_C > 0)$ | Tests whether there was a positive capital return during the period |
| $\mathbb{1}(R_C \leq 0)$ | Tests whether there was a negative capital return during the period |
| $\mathbb{1}(W_U - C_U > 0)$ | Tests whether untaxed wealth is greater than consumption from untaxed wealth |
| $\mathbb{1}(W_U - C_U \leq 0)$ | Tests whether consumption from untaxed wealth is greater than or equal to untaxed wealth |
| $\mathbb{1}(-L - L_U > W_T \times R_C \times TO_S)$ | Tests whether losses carried forward remaining after offsetting gains realized from consumption are greater than short-term losses incurred during the period |
| $\mathbb{1}(-L - L_U \leq W_T \times R_C \times TO_S)$ | Tests whether losses carried forward remaining after offsetting gains realized from consumption are less than or equal to short-term losses incurred during the period |
| $\mathbb{1}(-L - L_U > (W_U - C_U) \times TO_L)$ | Tests whether losses carried forward remaining after offsetting gains realized from consumption are greater than long-term gains realized from long-term turnover |
| $\mathbb{1}(-L - L_U \leq (W_U - C_U) \times TO_L)$ | Tests whether losses carried forward remaining after offsetting gains realized from consumption are less than or equal to long-term gains realized from long-term turnover |

| | |
|---|---|
| $\mathbb{1}(-L - L_U > ((W_U - C_U) \times TO_L) + ((W_T - C_A - C_U) \times R_C \times TO_S))$ | Tests whether losses carried forward remaining after offsetting gains realized from consumption are greater than long-term gains realized from long-term turnover and short-term gains realized from short-term turnover |
| $\mathbb{1}(-L - L_U \leq ((W_U - C_U) \times TO_L) + ((W_T - C_A - C_U) \times R_C \times TO_S))$ | Tests whether losses carried forward remaining after offsetting gains realized from consumption are less than or equal to long-term gains realized from long-term turnover and to short-term gains realized from short-term turnover |
| $\mathbb{1}(-L - L_U > (W_T - C_A - C_U) \times R_C \times TO_S)$ | Tests whether losses carried forward remaining after offsetting gains realized from consumption are greater than short-term gains realized from short-term turnover |
| $\mathbb{1}(-L - L_U \leq (W_T - C_A - C_U) \times R_C \times TO_S)$ | Tests whether losses carried forward remaining after offsetting gains realized from consumption are less than or equal to short-term gains realized from short-term turnover |
| $\mathbb{1}(C_U > -L_U)$ | Tests whether consumption from untaxed wealth is greater than losses available to offset consumption |
| $\mathbb{1}(C_U \leq -L_U)$ | Tests whether consumption from untaxed wealth is less than or equal to losses available to offset consumption |

**Consumption:**

$$cmf = \left[1 + \left(\frac{F}{(1+D)^{n_C-1}}\right)\right]^{n_C-1} \quad (A.1)$$

Implementing the consumption modification factor with the forward consumption rate (F) set to 3 percent and the forward consumption dampening rate ($D$) set to 2.75 percent modifies simple fractional wealth consumption so that it conforms to single life expectancy required minimum distributions (RMDs) from traditional IRAs as imposed by the U.S. Internal Revenue Service. Alternative $F$ and $D$ values should be tested and verified to meet RMDs.

$$\mathbf{C_A} = \mathbb{1}(n_A = 0) \times (W_A/n_C) \times cmf \quad (A.2)$$

If the investor is still in the accumulation period then consumption from after-tax wealth will be 0, otherwise it is equal to after-tax wealth divided by the expected number of years remaining in the consumption period multiplied by the consumption modification factor.

$$\mathbf{C_U} = \mathbb{1}(n_A = 0) \times (W_U/n_C) \times cmf \quad (A.3)$$

If the investor is still in the accumulation period then consumption from untaxed wealth will be 0, otherwise it is equal to unrealized gains/losses divided by the expected number of years remaining in the consumption period multiplied by the consumption modification factor.

**After-tax cash flow:**

From a Taxable Account

$$\mathbf{AC = C_A + C_U - T_N}$$

From a Tax-deferred account

$$\mathbf{AC = (C_A + C_U) \times (1 - T_M)}$$

From a tax-exempt account

$$\mathbf{AC = (C_A + C_U)}$$

**Total wealth:**

$$\mathbf{W_T = W_A + W_U}$$

$$\mathbf{(W_T)_{t+1} = (W_A)_{t+1} + (W_U)_{t+1}}$$

$$\mathbf{(W_A)_{t+1} =}$$

Sum of Equations (A.4) through (A.14)

$$(W_A - C_A) + ((W_T - C_A - C_U) \times R_I \times (1 - T_I)) \tag{A.4}$$

Recognize and tax income for the period.

$$+ (\mathbb{1}(W_U - C_U > 0) \times (W_U - C_U) \times TO_L \times (1 - T_L)) \tag{A.5}$$

If untaxed wealth exists then realize and tax long-term gains that result from long-term turnover.

$$+ (\mathbb{1}(W_U - C_\mathrm{U} \leq 0) \times (W_U - C_U) \times TO_L) \tag{A.6}$$

If losses exist then realize the long-term loss that results from long-term turnover.

$$+ (\mathbb{1}(W_U - C_U > 0) \times \mathbb{1}(-L - L_U > (W_U - C_U) \times TO_L) \times (W_U - C_U) \times TO_L \times T_L) \tag{A.7}$$

If untaxed wealth exists and a loss carry-forward exists that exceeds the realized long-term gain then use the loss carry-forward to fully offset the realized long-term gains and add back taxes deducted in Equation (A.5).

$$+ (\mathbb{1}(W_U - C_U > 0) \times \mathbb{1}(-L - L_U \leq (W_U - C_U) \times TO_L) \times (-L - L_U) \times T_L) \tag{A.8}$$

If untaxed wealth exists and a loss carry-forward exists that does not exceed the realized long-term gain then use the available loss carry-forward to partially offset the realized long-term gain and add back taxes deducted in Equation (A.5) to the extent that the available loss carry-forward offsets the gain.

$$+ (\mathbb{1}(R_C > 0) \times (W_T - C_A - C_U) \times R_C \times TO_S \times (1 - T_S)) \tag{A.9}$$

If the capital return for the period is positive then realize and tax short-term gains that result from short-term turnover.

$$+ (\mathbb{1}(R_C \leq 0) \times (W_T - C_A - C_U) \times R_C \times TO_S) \tag{A.10}$$

If the capital return for the period is negative then realize the short-term loss that results from short-term turnover.

$$+ (\mathbb{1}(R_C > 0) \times \mathbb{1}(W_U - C_U > 0) \times \mathbb{1}(-L - L_U > ((W_U - C_U) \times TO_L) + ((W_T - C_A - C_U) \times R_C \times TO_S)) \times (W_T - C_A - C_U) \times R_C \times TO_S \times T_S) \tag{A.11}$$

If the capital return for the period is positive, untaxed wealth exists, and a loss carry-forward exists that exceeds realized long-term and short-term gains then use the loss carry-forward to fully offset the realized short-term gain and add back taxes deducted in Equation (A.9).

$$+ (1(R_C > 0) \times \mathbb{1}(W_U - C_U > 0) \times \mathbb{1}(-L - L_U \leq ((W_U - C_U) \times TO_L) + ((W_T - C_A - C_U) \times R_C \times TO_S)) \times \mathbb{1}(-L - L_U > ((W_U - C_U) \times TO_L)) \times ((-L - L_U) - (W_U - C_U) \times TO_L) \times T_S) \tag{A.12}$$

If the capital return for the period is positive, untaxed wealth exists, and a loss carry-forward exists that does not exceed realized long-term and short-term gains but is greater than the realized long-term gain then use the loss carry-forward to partially offset the realized short-term gain and add back taxes deducted in Equation (A.9) to the extent the available loss carry-forward offsets the gain.

$$+ (\mathbb{1}(R_C > 0) \times \mathbb{1}(W_U - C_U \leq 0) \times \mathbb{1}(-L - L_U > (W_T - C_A - C_U) \times R_C \times TO_S) \times (W_T - C_A - C_U) \times R_C \times TO_S \times T_S) \tag{A.13}$$

If the capital return for the period is positive, untaxed wealth does not exist, and a loss carry-forward exists that exceeds the realized short-term gain then use the loss carry-forward to fully offset realized short-term gain and add back taxes deducted in Equation (A.9)

$$+ (\mathbb{1}(R_C > 0) \times \mathbb{1}(W_U - C_U \leq 0) \times \mathbb{1}(-L - L_U \leq (W_T - C_A - C_A) \times R_C \times TO_S) \times (-L - L_U) \times T_S) \quad \text{(A.14)}$$

If the capital return for the period is positive, untaxed wealth does not exist, and a loss carry-forward exists but does not exceed realized short-term gain then use the loss carry-forward to partially offset the realized short-term gain and add back taxes deducted in Equation (A.9) to the extent the available loss carry-forward offsets the realized short-term gain.

$$(\mathbf{L})_{\mathbf{t+1}} =$$

Sum of Equations (A.15) through (A.22).

$$((L + L_U) + (\mathbb{1}(W_U - C_U < 0) \times (W_U - C_U) \times TO_L)) \quad \text{(A.15)}$$

If unrealized losses exist then realize the long-term loss that results from long-term turnover.

$$+ (\mathbb{1}(W_U - C_U > 0) \times \mathbb{1}(-L - L_U > (W_U - C_U) \times TO_L) \times (W_U - C_U) \times TO_L) \quad \text{(A.16)}$$

If unrealized gains exist and the loss carry-forward exceeds realized long-term gains then reduce the loss carry-forward by the amount used to offset realized long-term gains. This reduction in carry-forward is equal to the carry-forward used in Equation (A.7).

$$+ (\mathbb{1}(W_U - C_U > 0) \times \mathbb{1}(-L - L_U \leq (W_U - C_U) \times TO_L) \times (-L - L_U)) \quad \text{(A.17)}$$

If unrealized gains exist and the loss carry-forward does not exceed realized long-term gains then reduce the loss carry-forward to the extent that the realized long-term gain is offset. This reduction in carry-forward is equal to the carry-forward used in Equation (A.8).

$$+ (\mathbb{1}(R_C \leq 0) \times (W_T - C_A - C_U) \times R_C \times TO_S) \quad \text{(A.18)}$$

If the capital return for the period is negative then increase the loss carry-forward by the short-term loss realized as a result of short-term turnover.

$$+ (\mathbb{1}(R_C > 0) \times \mathbb{1}(W_U - C_U > 0) \times \mathbb{1}(-L - L_U > ((W_U - C_U) \times TO_L) + ((W_T - C_A - C_U) \times R_C \times TO_S)) \times (W_T - C_A - C_U) \times R_C \times TO_S) \quad \text{(A.19)}$$

If the capital return for the period is positive, untaxed wealth exists, and a capital loss carry-forward exists that exceeds the realized long-term and short-term gains then reduce the loss carry-forward by the realized short-term gain it offset. This reduction in carry-forward is equal to the carry-forward used in Equation (A.11).

$$+ (\mathbb{1}(R_C > 0) \times \mathbb{1}(W_U - C_U > 0) \times \mathbb{1}(-L - L_U \leq ((W_U - C_U) \times TO_L) + ((W_T - C_A - C_U) \times R_C \times TO_S)) \times \mathbb{1}(-L - L_U > (W_U - C_U) \times TO_L) \times ((-L - L_U) - ((W_U - C_U) \times TO_L))) \quad \text{(A.20)}$$

If the capital return for the period is positive, untaxed wealth exists, and a loss carry-forward exists that does not exceed the realized long-term and short-term gains but is greater than the realized long-term gain then reduce the capital loss carry-forward to the extent that the carry-forward exceeds the realized long-term gain and offsets the realized short-term gain. This reduction in carry-forward is equal to the carry-forward used in Equation (A.12).

$$+ (\mathbb{1}(R_C > 0) \times \mathbb{1}(W_U - C_U \leq 0) \times \mathbb{1}(-L - L_U > (W_T - C_A - C_U) \times R_C \times TO_S) \times (W_T - C_A - C_U) \times R_C \times TO_S) \quad (A.21)$$

If the capital return for the period is positive, untaxed wealth does not exist, and a loss carry-forward exists which exceeds the realized short-term gain then reduce the loss carry-forward by the amount used to offset the realized short-term gain. This reduction in carry-forward is equal to the carry-forward used in Equation (A.13).

$$+ (\mathbb{1}(R_C > 0) \times \mathbb{1}(W_U - C_U \leq 0) \times \mathbb{1}(-L - L_U \leq (W_T - C_A - C_U) \times R_C \times TO_S) \times (-L - L_U)) \quad (A.22)$$

If the capital return for the period is positive, untaxed wealth does not exist, and a capital loss carry-forward exists but does not exceed the realized short-term gain then reduce the loss carry-forward to the extent the realized short-term gain is offset. This reduction in carry-forward is equal to the carry-forward used in Equation (A.14).

$$\mathbf{(W_U)_{t+1}} =$$

Sum of Equations (A.23) and (A.24).

$$(W_U - C_U) \times (1 - TO_L) \quad (A.23)$$

Reduce untaxed wealth by the long-term gain realized in Equation (A.5) as a result of long-term turnover.

$$+(W_T - C_A - C_U) \times R_C \times (1 - TO_S) \quad (A.24)$$

Increase untaxed wealth by the capital return experienced during the period less the short-term gain realized during the period as a result of short-term turnover.

$$\mathbf{L_U} =$$

Sum of Equations (A.25) and (A.26).

$$\mathbb{1}(C_U > -L) \times -L \quad (A.25)$$

If consumption from untaxed wealth is greater than the loss carry-forward available then use the loss carry-forward to the extent it is available to offset gains realized as a result of consumption.

$$+\mathbb{1}(C_U \leq -L) \times C_U \quad (A.26)$$

If consumption from untaxed wealth is less than the loss carry-forward available then completely offset gains realized as a result of consumption. This will also realize losses from consumption of untaxed wealth if an unrealized loss exists.

$$\mathbf{T_N} = (C_U \times T_L) - (L_U \times T_L) \quad (A.27)$$

The net tax liability is equal to the tax liability from realizing gains as a result of consumption less the offset provided by using loss carry-forwards.

## Notes

1 The term tax-exempted is used throughout this paper to describe variables which have not been adjusted for taxes. For standard tax-exempted asset allocation inputs, it can be said that the analyst conducting the portfolio analysis has either implicitly or explicitly exempted the inputs from the impact of taxes.

2 Subtracting income returns from total returns to determine returns from capital appreciation requires addressing the possibility of extreme negative simulated total returns. Consider a simulated negative total return of 99 percent and a simulated positive income return of 4 percent. This would imply a capital loss of 103 percent in order for the sum of income and capital gain to equal the simulated total return. To address such instances simulated income returns were adjusted when necessary to prevent the sum of the simulated total return and the income return from exceeding a loss of greater than 100 percent. In the example provided the simulated income return would be adjusted to equal 1 percent.

3 See Internal Revenue Code § 1.401(a)(9) or IRS Publication 590.

4 Modified fractional consumption is determined as follows:

$$C = \frac{W_T}{n_C} \times cmf$$

Where:

$W_T =$ Total wealth

$n_C =$ Years remaining in consumption period

And:

$$cmf = Consumption\ modification\ factor = \left[1 + \left(\frac{F}{(1+D)^{n_C-1}}\right)\right]^{n_C-1}$$

$F =$ Forward consumption rate

$D =$ Forward consumption dampening rate

5 Transformation of lognormal expected present values were accomplished using the following equations:

$$E(PV) = \exp\left\{\mu_{\ln PV} + \frac{1}{2}\sigma^2_{\ln PV}\right\}$$

$$\sigma^2_{PV} = E(PV^2) - E(PV)^2$$

$$\mathrm{Cov}_{PV_1,PV_2} = E(PV_1 \times PV_2) - [E(PV_1) \times E(PV_2)]$$

Where:

$$E(PV^2) = \exp\{2\mu_{\ln PV} + 2\sigma^2_{\ln PV}\}$$

$$E(PV_1 \times PV_2) = \exp\left\{\mu_{\ln(PV_1 \times PV_2)} + \frac{1}{2}\sigma^2_{\ln(PV_1 \times PV_2)}\right\}$$

$$\mu_{\ln(PV_1 \times PV_2)} = E(\ln PV_1) + E(\ln PV_2)$$

$$\sigma^2_{\ln(PV_1 \times PV_2)} = \sigma^2_{PV_1} + \sigma^2_{PV_2} + 2Cov_{\ln PV_1, \ln PV_2}$$

6 Normal (Gaussian) return distributions are sufficient to justify a mean–variance analysis, but they are not necessary (see Markowitz and Blay, 2013). In this instance a mean–variance analysis is conducted using investments with lognormally distributed present value outcomes.

7 Inputs were derived using common input estimation methods and are for illustrative purposes only.

8 The tax rates used are for illustrative purposes and may not coincide with the current U.S. Federal Tax Code at the time of publication.

9 Our analysis assumes an asset class implementation using ETF index investments. The turnover rates applied assume that the ETFs are expected to incur negligible capital gain distributions.

10 A futures-based commodity implementation is assumed in our analysis. Gains for these types of investments are taxed annually at a blended rate comprised of 60% long-term capital gains tax rate and 40% short-term capital gains rate regardless of whether the investment was sold or not.

11 The mean–variance optimizer used was developed by G. Peter Todd PhD, CFA of Parilux Investment Technology, LLC.

12 The conversion of efficient frontiers based on tax-exempted inputs into present value terms was accomplished by converting individual frontier portfolios using the following equations:

$$E(PV_p) = [PV\mu_1 \quad PV\mu_2 \quad \cdots \quad PV\mu_n] \times \begin{bmatrix} w_1 \\ w_2 \\ \vdots \\ w_n \end{bmatrix}.$$

$$\sigma^2_p = [w_1 \quad w_2 \quad \cdots \quad w_n] \times \begin{bmatrix} \sigma_{1,1} & \sigma_{1,2} & \cdots & \sigma_{1,n} \\ \vdots & \vdots & & \vdots \\ \sigma_{n,1} & \sigma_{n,2} & \cdots & \sigma_{n,n} \end{bmatrix} \times \begin{bmatrix} w_1 \\ w_2 \\ \vdots \\ w_n \end{bmatrix}.$$

Where:

$PV\mu_j =$ Expected present value for asset $j$

$w_j =$ The weight of asset $j$ in the portfolio

$\sigma_{i,j} =$ Covariance of the present values ofasset $i$ with the present values ofasset $j$.

13 For a more objective comparison the optimizer was allowed to allocate to municipal bonds in the taxable account.

14 We have presented TCPA and MVO frontiers that have not been constrained beyond the limitations imposed by the availability of investor assets in different account types. In practice, the application of additional constraints to limit asset class concentrations within portfolios may be prudent.

15 The average annual real after-tax cash flow is calculated as follows:

$$\overline{RACF} = d_A \times \left[\frac{\sum_{r=1}^{R} c_r \times d_r}{\sum_{r=1}^{R} d_r}\right]$$

Where:

$d_A =$ discount factor for accumulation period $= \frac{1}{(1+s)^A}$

$d_r =$ discount factor for consumption cash flow at time $r = \frac{1}{(1+s)^{r-1}}$

and

$s =$ intertemporal rate of substitution

$A =$ Total number of years in accumulation

$R =$ Total number of years in consumption

$r =$ consumption year

$c_r =$ After-tax consumption cash flow at time $r$

16 Left tail lnPV values for specific confidence levels are identified using the following equation:

$$lnPV_C = \mu_{ln\,PV} + (-z \times \sigma_{ln\,PV})$$

Where:

$ln\,PV_C =$ lnPV at confidence level $C$

$\mu_{ln\,PV} =$ Mean of the lnPV distribution

$z =$ The $z$-value corresponding to the specific confidence level used

$\sigma_{ln\,PV} =$ Standard deviation of the lnPV distribution.

17 The weighted average present value outcomes of a portfolio of investments that are individually lognormally distributed will generally not be exactly lognormally distributed. The assumption that portfolio outcomes are lognormal allows for the determination of approximate present value outcomes.

18 For resampled portfolios it is typically not true that the average of the sample portfolio standard deviations is equal to the standard deviation of the average (resampled) portfolio. Therefore, we calculate portfolio means and standard deviations using constituent means, standard deviations, and correlations using large samples of present values. These large samples, which are conveniently produced by aggregating simulations required for the resampling process, provide a reasonable basis for the estimation of portfolio confidence level outcomes.

## References


Blanchett, D. and Kaplan, P. (2013). "Alpha, Beta, and Now. . .Gamma," *The Journal of Retirement* **1**(2), 29–45.

Dammon, R. M., Spatt, C. S., and Zhang, H. H. (2004). "Optimal Asset Location and Allocation with Taxable and Tax-Deferred Investing," *Journal of Finance* **59**(3) (June), 999–1037.

Horan, S. M. (2007). "Applying After-Tax Asset Allocation," *The Journal of Wealth Management* **10**(2), 84–93.

Horan, S. M. and Al Zaman, A. (2008). "Tax-Adjusted Portfolio Optimization and Asset Location: Extensions and Synthesis," *The Journal of Wealth Management* **11**(3), 56–73.

Horan, S. M. and Adler, D. (2009). "Tax Aware Investment Management Practice," *The Journal of Wealth Management* **12**(2), 71–88.

Jeffrey, R. H. and Arnott, R. D. (1993). "Is Your Alpha Big Enough to Cover Its Taxes?" *The Journal of Portfolio Management* **19**(3), 15–25.

Jennings, W. W., Horan, S. M., Reichenstein, W., and Brunel, L. P. (2011). "Perspectives from the Literature of Private Wealth Management," *The Journal of Wealth Management* **12**(2), 71–88.

Markowitz, H. M. (1959). *Portfolio Selection: Efficient Diversification of Investments*, second edition. Malden, Massachusetts: Blackwell Publishers.

Markowitz, H. M. and Blay, K. A. (2013). *Risk-Return Analysis: The Theory and Practice of Rational Investing*, first edition. New York, New York: McGraw-Hill.

Reichenstein, W. (2000). "Calculating the Asset Allocation," *Journal of Wealth Management* **3**(2), 20–25.

Reichenstein, W. (2001). "Asset Allocation and Asset Location Decisions Revisited," *Journal of Wealth Management* **4**(1), 16–26.

Reichenstein, W. (2004). "Tax-Aware Investing: Implications for Asset Allocation, Asset Location, and Stock Management Style," *Journal of Wealth Management* **7**(3) (January), 7–18.

Reichenstein, W. (2006). "After-Tax Asset Allocation," *Financial Analysts Journal* **62**(4), 14–19.

Reichenstein, W. (2007). "Implications of Principal, Risk, and Return Sharing across Savings Vehicles," *Financial Services Review* **16**(1), 1–17.

Reichenstein, W. (2007). "Note on Applying After-Tax Asset Allocation," *Journal of Wealth Management* **10**(2), 94–97.

Shoven, J. B. and Sialm, C. (2003). "Asset Location in Tax-Deferred and Conventional Savings Accounts," *Journal of Public Economics* **88**, 23–38.

Stuart, A. and Ord, K. (2004). *Kendall's Advanced Theory of Statistics*, West Sussex, UK: John Wiley & Sons Ltd.

Wilcox, J., Horvitz, J. E., and diBartolomeo, D. (2006). *Investment Management for Taxable Private Investors*. Charlottesville, VA: CFA Institute Publications.


*Keywords*: Asset allocation; asset location; tax efficiency; tax-aware asset allocation; tax-efficient asset allocation; mean–variance optimization of present values