FAJ Financial Analysts Journal
Volume 71 • Number 4

PERSPECTIVES

What Would Yale Do If It Were Taxable?

Patrick Geddes, Lisa R. Goldberg, and Stephen W. Bianchi, CFA

The distinctive financial goals and constraints of ultra-high-net-worth individuals together with their aggregate growth in assets have led to the emergence of "New Institutional" investing, which includes the best practices from institutional investors but also incorporates the critical element of tax management. The authors design New Institutional asset allocations that incorporate traditional investment metrics in a tax-aware setting. Specifically, they show how risk and after-tax returns need to be combined from inception when seeking an optimal after-tax asset allocation. Diversification is especially important for taxable investors because low asset class correlations can facilitate the inclusion of attractive but tax-inefficient asset classes in a tax-aware allocation.

Asset allocation is an essential element of every investment process, and much thought has been devoted to the subject. Historically, most of the literature has revolved around institutional investors because they have accounted for most of the growth in assets under management.

Since the start of the millennium, however, and especially in the wake of the financial crisis, a "New Institutional" class of ultra-high-net-worth (UHNW) investors has emerged. Asset allocations for new and traditional institutional investors differ because UHNW investors pay taxes whereas most institutional investors do not. Taxes make the wholesale adoption of a successful pretax allocation a losing strategy for New Institutional investors. Constructing a winning strategy requires additional analysis along many dimensions.[1] In this article, we show how tax management and risk control are inseparable elements of asset allocation for a taxable investor.

The differences between a successful pretax allocation and its after-tax counterpart are fundamental. We illustrate the differences using Yale University's endowment as a baseline.[2] The Yale asset allocation is predicated on diversification and also on the premium for illiquidity that can be wrung from private asset classes by investors who are either skilled or lucky.[3] These features are just as attractive to a taxable investor as they are to a tax-exempt investor. We show how to adapt the Yale allocation (or any other pretax asset allocation) to make it more tax efficient while preserving its desirable attributes.

Patrick Geddes is CEO at Aperio Group, Sausalito, California. Lisa R. Goldberg is director of research at Aperio Group and adjunct professor of economics and statistics at the University of California, Berkeley. Stephen W. Bianchi, CFA, is a consultant at Aperio Group and affiliated researcher at the Center for Risk Management Research, University of California, Berkeley.

Some Guidelines for After-Tax Asset Allocation

The complexity of tax law, its variation over time, and its wide range of effects on different portfolios present a challenge to even the most tax-aware adviser. These considerations complicate the task of identifying after-tax investment principles that are universally applicable. Nonetheless, it is important to not let the perfect be the enemy of the good when assessing the level of precision necessary to estimate tax impact. Some general tax effects are quite simple. For example, equity index funds tend to be more tax efficient than actively managed equity funds.[4]

In our study, we considered seven asset classes and made specific assumptions about the tax efficiency of each one. Inevitably, our results depend on our assumptions, so our examples are illustrations that may not reflect the situation of any particular investor or time period. Nevertheless, some high-level observations emerge. For example, our analysis shows ways to offset the tax drag that accrues in actively managed equity funds and hedge funds, allowing them to be included in an after-tax asset allocation in spite of their tax inefficiency.[5]

To create an optimal after-tax asset allocation, we began with an optimal pretax asset allocation and derived its implied returns. Subsequently, we adjusted those pretax returns for taxes and then reoptimized.[6] This analysis shows profound differences

between the pretax and after-tax allocations. The result is an investment strategy that takes advantage of the diversification benefits built into a pretax strategy while addressing the tax realities of the most affluent individual investors.[7] New Institutional investors will increasingly expect such rigorous tax analysis as part of what they pay their advisers for.

Ferreting Out Expected Returns

In order to modify a pretax allocation for a taxable investor, we needed to assess the impact of taxes on pretax expected returns. Of course, the task of forecasting pretax returns is famously difficult,[8] with advisers relying on tools ranging from historical estimation to macroeconomic forecasts. We skirted this issue by using *reverse optimization*, based on a publicly available asset allocation and a covariance matrix,[9] to determine expected returns. This approach was first proposed by Sharpe (1974).

A standard mean–variance optimization takes expected returns as inputs and produces optimal portfolio weights. If we know the weights but not the expected returns, however, we simply reverse the standard process to solve for a different output, as shown in **Figure 1**.

Taking the publicly available asset class weights of Yale's endowment as a starting point,[10] we backed in to the endowment's expected returns with reverse optimization. The covariance matrix was estimated from representative benchmark indexes from the Morningstar Principia database.[11]

Table 1 shows the Yale asset class weights as of 31 December 2013, with foreign and US equities collapsed into one asset class, along with the implied asset class returns from the reverse optimization.

These implied returns may come as a shock—especially the low 2.2% implied return for absolute return (hedge funds). This implied return does not suggest that the consensus return expectation for hedge funds is only 2.2% or that Yale's investment experts expect that exact return. It simply implies that, on the basis of the historical covariance among those assets, a hedge fund return of 2.2% is the only return that is mathematically consistent with a reverse-optimized mean–variance analysis using the weights of the Yale endowment. The weights and covariance can produce no other answer. The expected return for hedge funds could be higher only if the Yale weight were higher or the historical covariance matrix were different. Similarly, a change in implied return for any of the other asset classes would require a change in the inputs to the reverse optimization. To put the matter in perspective, consider that a much higher expected return assumption in combination with the low volatility of hedge funds and their presumed low correlation with other asset classes could be grounds for the extreme case of allocating 100% of the portfolio to hedge funds.

In Appendix A, we show the indexes used as proxies for the asset class returns in the estimation of our historical covariance matrix. A different set of

Figure 1. Two Applications of a Mean–Variance Optimizer

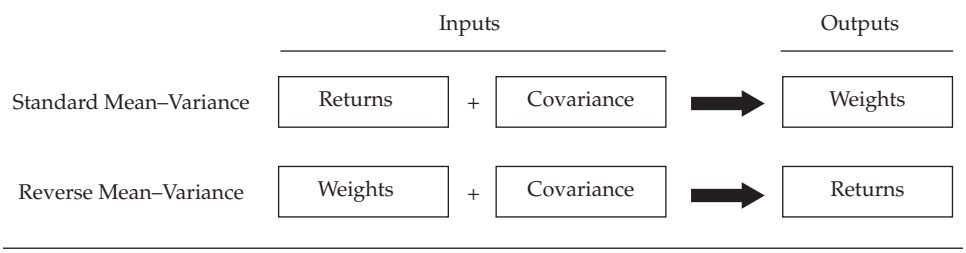

Table 1. Asset Class Weights and Implied Returns for the Yale Endowment

Asset Class	Benchmark	Pretax Weight	Implied Pretax Return
Absolute return	Credit Suisse AllHedge Index	17.8%	2.2%
World public equity	MSCI ACWI Index	15.7	9.7
Bonds	Barclays U.S. Aggregate Bond Index	4.9	1.5
Natural resources	S&P North American Natural Resources Index	7.9	10.7
Real estate	Dow Jones U.S. Real Estate Index	20.2	12.3
Private equity	Russell 2000 Index	32.0	11.7
Cash	BofA Merrill Lynch 0–3 Month US Treasury Bill Index	1.5	1.5

Notes: Past performance is not a guarantee of future returns. The performance shown in the table is hypothetical and does not represent the actual performance of the Yale endowment or any other portfolio.

Sources: Weights are from the annual report of the Yale Endowment as of 31 December 2013; returns were calculated by Aperio Group using the reverse-optimization model described herein.

proxies for these asset classes, a different time period, or a different exponential weight for the estimation of the covariance matrix would produce different results that could be equally valid.[12] For example, Table 1 reflects a hedge fund benchmark with a very low correlation with equity markets, which is one reason that the implied expected return is so low for the absolute return asset class. Later in this article, we will look at the impact of choosing a different hedge fund benchmark.

Investors who seek to modify specific asset class return forecasts to reflect their own expectations can use the well-known model developed by Black and Litterman (1992), which expanded on Sharpe's original model. In the Black–Litterman model, investors can adjust expected returns up or down in order to incorporate their views.

The Painful Yet Often-Ignored Impact of Taxes

The next step is to adjust the pretax returns to reflect the tax impact on various asset classes. The tax impact, or "tax haircut," converts pretax total return into after-tax total return. Pretax return can be divided into two components—the unrealized portion and the realized portion, which can include ordinary income, dividends, and realized short- and long-term gains.

The higher the proportion of unrealized return, the more effectively the taxable investor can control the timing of the haircut.[13] **Figure 2** shows how for a given pretax return, an investment with a larger proportion of realized gain will lead to a higher tax burden (and thus a lower after-tax return) than a more tax-efficient investment, where less of the pretax return is subject to the tax penalty.

The magnitude of the tax haircut varies greatly across asset classes, of course, but it also can vary a great deal within one specific asset class. Following are descriptions of the issues faced in estimating the tax haircut for major asset classes.

1. *Public equities.* In the public equity space, an investor can choose three different investment approaches: (1) traditional active management, (2) traditional indexing, or (3) tax-advantaged indexing, also known as tax-loss harvesting. For equity strategies, taxable income can come from either dividends or the realization of gains by an investor or the investor's manager. Active management, on average, has tended to incur the worst tax haircut. As a historical example, we gathered data for all active US equity funds from the Morningstar Principia database over the 18-year period ending 30 June 2013. We assumed that the dividend distribution would be the same for all three types of public equity strategies (i.e., all would pay dividends comparable to those of the entire stock market). If all three paid similar dividends, that would leave only the breakdown between realized and unrealized capital gains. Over this period, 76% of the total capital appreciation from actively managed funds was distributed as taxable income to investors. That percentage varies over time, of course, but here we treated it as representative of the tax haircut of active management, rounded the realized portion down to 70% per year, and applied the same haircut to global equities.

Figure 2. Breakdown of Realized and Unrealized Return

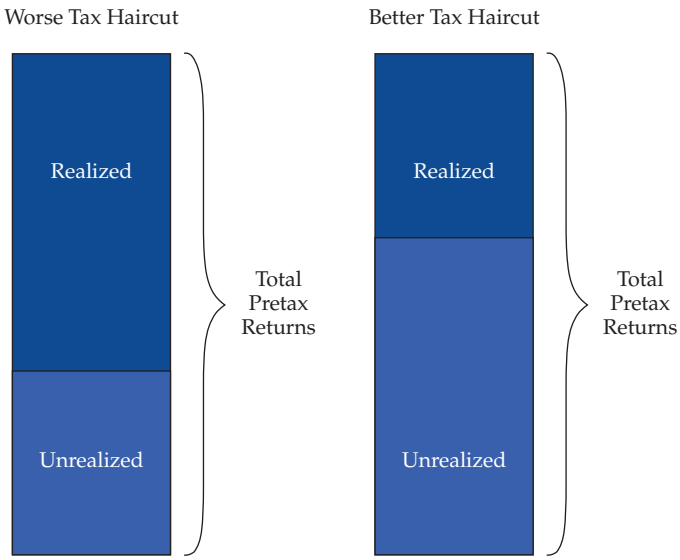

Note: This figure is for illustrative purposes only.

Turning from active management to traditional indexing, although the dividends are presumed to remain the same, the realized gain component has tended to be much smaller (even zero). In our analysis, we assumed that it was 0%, reflecting the fact that for the many efficiently managed and broadly diversified large index funds, capital gains can be kept very low. Furthermore, many exchange-traded funds have taken advantage of the tax rules concerning redemptions of units to avoid capital gain distributions.

Finally, investors can choose a third type of equity strategy known as tax-advantaged indexing, or tax-loss harvesting, which seeks to track a benchmark before tax—albeit with higher tracking error than a pure index fund—while generating a tax benefit from realizing losses. Such a strategy can actually go further in improving the tax haircut than index funds, which by law cannot pass through losses. Tax-advantaged indexing can generate a positive impact on after-tax returns, assuming that a taxable investor has realized gains elsewhere that can be offset with losses. For a federal-only taxable investor in the top tax bracket, that benefit has translated into a gain as high as 1.9% per year for assets passed through an estate or to a charity. The benefit has been as high as 1.0% for investors planning to liquidate at the end of a 10-year holding period.[14] Such gains from what is effectively tax deferral can be earned over longer periods of time—as long as 20 or 30 years—even when loss-harvesting opportunities diminish in later years.[15] In states with high tax rates, the benefit is magnified.[16]

2. *Bonds and cash.* It is straightforward to estimate tax haircuts for bonds and cash, which we assumed to have 100% of total return distributed as taxable ordinary income, with no capital gains. When the Yale portfolio is analyzed as tax exempt, the bond benchmark used as a proxy in the asset allocation is the taxable Barclays U.S. Aggregate Bond Index; when the Yale portfolio is treated as taxable, the bond benchmark is the tax-exempt Barclays U.S. Municipal Index, with all the returns treated as tax exempt.
3. *Absolute return.* Estimating a tax haircut for such strategies as hedge funds presents a challenge because public after-tax data are not readily available. On the basis of an informal survey of a select group of UHNW wealth managers, in our analysis, we assumed the same haircut on gains for hedge funds as for actively managed equities.
4. *Natural resources.* Investments in natural resources also present the problem of very limited public data with respect to the magnitude of the tax haircut. Furthermore, the method of capturing exposure to natural resources can significantly affect the tax haircut. Futures-based natural resource strategies may face significant tax penalties, whereas direct ownership can be highly tax efficient owing to the preferential tax treatment of certain investments in timber or energy, for example.[17] For this analysis, we assumed that 30% of gains for natural resources are realized, with no ordinary income. In reality, investors face a wide range of tax treatments for natural resource investments, but they were defined in this analysis as relatively tax efficient, certainly compared with active equity or hedge funds.
5. *Real estate.* Similar to natural resources, real estate investing offers structures with a wide range of tax treatment. Direct investment or pass-through pooled vehicles, such as partnerships, can offer highly advantageous tax treatment. Ownership through real estate investment trusts (REITs), however, does not offer the same tax benefits. In fact, they can be tax inefficient because REIT dividends do not qualify for the beneficial dividend tax rate; instead, they are subject to the highest rate for ordinary income. For this analysis, we assumed that 30% of total return is ordinary income, with no realized gains.
6. *Private equity.* Although private equity can offer tax deferral at capital gains rates, many investors capture this asset class through funds, which *do* distribute gains and other income to investors. Again, we applied the same relatively tax-efficient assumptions, with 30% of gains realized each year, all of it long term. Dividend assumptions are the same as those for public equities.

All the assumptions about taxation by asset class can be seen in **Table 2**. Long-term capital gains are assumed to constitute two-thirds of total realized capital gains, with the remaining one-third being treated as short term. An exception is made for private equity, where all gains are assumed to be long term.[18]

The tax haircut estimates in Table 2 reflect disposition through charity or an estate, where a step-up in basis would eliminate the effect of taxation on liquidation.[19] Estate taxes could further erode after-tax wealth, but in this article, we will focus exclusively on the effects of income taxes. The assumptions disregard the effect of taxes on liquidation, which could reduce the distinction among asset classes.

These tax haircut estimates represent just one of many valid sets of assumptions for a top-bracket taxpayer. The goal here is not to derive the unattainable "true" set of tax assumptions but, rather, to focus on the transparency of measuring and estimating actual

Table 2. Estimation of Tax Penalty across Asset Classes

Asset Class	Percentage of Total Return from Ordinary Income	Realized Short Gains	Realized Long Gains	Unrealized Capital Gains	Dividend Return	Tax Benefit
Absolute return	0%	23%	47%	30%	0.0%	0.0%
Active equity[a]	0	23	47	30	2.0	0.0
Indexed equity[a]	0	0	0	100	2.0	0.0
Tax-advantaged equity[a]	0	0	0	100	2.0	1.9
Taxable bonds	100	0	0	0	0.0	0.0
Municipal bonds	0	0	0	0	0.0	0.0
Natural resources	0	10	20	70	0.0	0.0
Real estate	30	0	0	70	0.0	0.0
Private equity[a]	0	0	30	70	2.0	0.0
Cash	100	0	0	0	0.0	0.0

Note: The estimates in this table are hypothetical and are not based on specific individual investments.

[a]For equity strategies, the percentage of total return applies only to the nondividend portion of total return (i.e., gains). For all public and private equity strategies, dividends are assumed to be taxed at qualified dividend rates applied to a constant 2.0% dividend yield.

tax impact instead of dismissing it as too complicated or irrelevant. With tax rates as high as 52%,[20] the tax haircut can potentially wipe out more than half the pretax return, a reduction of wealth that can dwarf other costs.

How a Great Pretax Asset Allocation Can Look So Much Worse after Taxes

Our analysis consists of the following three steps:

1. Using reverse optimization, calculate the implied pretax returns for a portfolio given the existing asset allocation.
   - Inputs: pretax weights and estimated correlations and volatilities
   - Output: expected pretax returns
2. Apply the tax haircuts to adjust the pretax returns to reflect the actual after-tax experience of a top-bracket taxable investor.
   - Input: expected pretax returns
   - Output: expected after-tax returns
3. Use mean–variance optimization to generate an ideal set of weights for the same asset classes but for a taxable investor.
   - Inputs: expected after-tax returns and pretax correlations and volatilities
   - Output: after-tax weights

Table 3 shows implied returns, both before and after taxes, and adjusted weights (once the Yale portfolio becomes subject to taxes). The variations reflect three different assumptions about how the portfolio includes exposure to public equities. As one can see from the table, it turns out that the choice of strategy for public equities can lead to very different overall asset allocations.

How would the asset allocation of Yale's endowment need to change if it were taxable? The answer depends on how public equity as an asset class is implemented in the portfolio. Following are detailed explanations of each version of after-tax portfolio implementation.

1. The After-Tax Active Equity column shows the weights of the Yale portfolio with the equity allocation available only through active strategies—thus subject to realizing a tax haircut of 70% of total capital gains each year. In this setting, active equity and absolute return (hedge fund) allocations become far less attractive because of the significant tax haircut crimping what had been more attractive pretax returns. Municipal bonds earn a high allocation owing to lower risk, whereas the large weight in private equity reflects its relative tax efficiency. Private equity, reflecting Yale's choice of the Russell 2000 Index as the benchmark, provides the least painful method for achieving equity exposure once the tax drag is applied.
2. The After-Tax Indexed Equity column shows the portfolio under the same assumptions as in Scenario 1 but with indexed equities now available as well. Given that the tax haircut on capital gains does not apply to this asset class, the portfolio composition changes significantly, with equities and bonds weighted heavily and a smaller portion in private equity.
3. The After-Tax Tax-Advantaged Equity column shows the portfolio with access to the same indexed strategy but with the inclusion of tax-advantaged equities as well. The allocation to the tax-advantaged equity strategy is limited by the amount of capital gains generated by other asset

Table 3. Implied Returns and Weights for Yale Endowment (Low-Correlation Hedge Fund Index)

	Returns			Weights			
Asset Class	Pretax Implied Return	After-Tax Implied Return	Tax Haircut	Pretax	After-Tax Active Equity	After-Tax Indexed Equity	After-Tax Tax-Advantaged Equity
Absolute return	2.2%	1.7%	–0.5%	17.8%	0.0%	0.0%	11.9%
Active equity	9.7	7.5	–2.2	15.7	0.0	0.0	0.0
Indexed equity	9.7	9.2	–0.5	0.0	0.0	45.6	0.0
Tax-advantaged equity	9.7	11.1	1.4	0.0	0.0	0.0	26.4
Bonds	1.5	0.9	–0.7	4.9	35.0	25.8	14.5
Natural resources	10.7	9.7	–1.0	7.9	12.2	0.0	5.2
Real estate	12.3	10.7	–1.6	20.2	13.9	9.0	2.8
Private equity	11.7	10.4	–1.3	32.0	38.9	19.6	39.1
Cash	1.5	0.8	–0.7	1.5	0.0	0.0	0.0

Notes: Bonds are assumed to be taxable for the pretax allocation and tax-exempt municipals for the after-tax portfolios. All distributions from tax-exempt bonds are assumed to be tax exempt (i.e., possible capital gains are ignored). The weights and returns in this table are hypothetical and are not based on actual investments. Cash has been constrained in the optimization to show a weight of zero.

classes, allowing only enough to match what can be offset directly within the portfolio each year. In this scenario, hedge funds, which had not been attractive when only active or standard indexing strategies were available, are attractive owing to the capacity of the tax-advantaged strategy to offset the presumed tax inefficiency of hedge funds.

What if we look beyond the Yale endowment? Appendix B shows results for the average of a broad range of university endowments, based on data from the National Association of College and University Business Officers (NACUBO). The NACUBO data show a lower allocation to less liquid asset classes, such as private equity, which is not surprising given the Yale endowment's size and emphasis on capturing a perceived illiquidity premium. Besides that difference, adjusting for taxes for NACUBO provides similar, and intuitive, results: the more tax-inefficient asset classes no longer provide the same value once an investor has to pay taxes.

In the preceding examples, we used university endowment allocations as pretax starting points for our analysis. Now we create a pretax starting point based on dialogues with UHNW advisers. The sample portfolio analyzed in **Table 4** represents just one plausible allocation and does not reflect any empirical data.

As with the Yale portfolio, active equity carries such a high tax penalty that the optimal portfolio avoids it entirely, reallocating to hedge funds and other alternatives. Hedge funds still receive some allocation, but much less than in the pretax version, except in the case where tax-advantaged equity can offset part of the tax penalty. The value of hedge funds, however, still depends very much on their correlation with equity markets, as we will see in the next section.

Table 4. Implied Returns and Weights for Sample UHNW Adviser (Low-Correlation Hedge Fund Index)

	Returns			Weights			
Asset Class	Pretax Implied Return	After-Tax Implied Return	Tax Haircut	Pretax	After-Tax Active Equity	After-Tax Indexed Equity	After-Tax Tax-Advantaged Equity
Absolute return	2.5%	1.9%	–0.5%	25.0%	13.8%	11.3%	24.5%
Active equity	10.9	8.4	–2.5	30.0	0.0	0.0	4.4
Indexed equity	10.9	10.4	–0.5	0.0	0.0	51.0	0.0
Tax-advantaged equity	10.9	12.3	1.4	0.0	0.0	0.0	24.6
Bonds	1.7	0.9	–0.7	20.0	42.3	34.1	23.2
Natural resources	12.0	10.8	–1.1	5.0	13.1	0.0	5.7
Real estate	12.9	11.2	–1.7	10.0	8.0	2.5	0.0
Private equity	12.1	10.8	–1.3	10.0	22.9	1.1	17.5
Cash	1.5	0.8	–0.7	0.0	0.0	0.0	0.0

Note: See notes to Table 3.

The Interplay between Taxes and Diversification

So far, our results have shown the changes to asset class weights when a tax-exempt portfolio becomes subject to taxation. Beyond just the tax haircut, though, the correlations between asset classes affect the outcome significantly. For example, the correlation between equities and hedge funds (absolute return) depends materially on the chosen hedge fund proxy. In the previous section, we used the Credit Suisse AllHedge Index, which had a low correlation with equities over the sample period, 31 December 1998 through 30 June 2013. Other hedge fund indexes' returns have shown a much higher correlation with equity returns over the same period; an example is shown in **Table 5**.

The hedge fund correlation with equities materially changes the outcome. We saw in Table 3 how hedge funds disappear from the allocation once taxes are introduced and reappear only when the constrained tax-advantaged strategy allows the portfolio to gain from the potential risk-reducing benefits of hedge funds enough to offset the tax haircut. Thus, when equity strategies are active or traditionally indexed in the model, the tax penalty on hedge funds outweighs any correlation benefit unless the tax effect is mitigated by the gains from tax-advantaged equity. However, if a hedge fund strategy reflects the risk patterns of the HFRI Fund Weighted Composite Index, with its higher correlation with equities, then the optimizer never allocates anything to hedge funds, no matter how public equities are implemented, as shown in **Table 6**.

Note that the expected pretax return for absolute return strategies (hedge funds) is higher than that based on the low-correlation benchmark shown in Table 3. This outcome reflects the fact that in reverse optimization, higher implied return to an asset class results from higher allocation, higher volatility, or higher correlation. Even though its implied return is higher, however, an allocation to highly correlated hedge funds still cannot be justified in the presence of taxes. When presented with this observation, one wealth manager commented, "Well of course I wouldn't pay the tax penalty to be in a tax-inefficient hedge fund strategy unless I were gaining the risk advantages of low correlation."

Table 5. Hedge Fund Index Correlation with Equities

Index Benchmark for Absolute Return (Hedge Funds)	Correlation with Equities[a]
Credit Suisse AllHedge Index (low correlation)	0.14
HFRI Fund Weighted Composite Index (high correlation)	0.88

[a]As measured by the MSCI ACWI Index.

After-Tax Allocation with a Tax-Adjusted Covariance Matrix

After-tax returns of tax-inefficient assets tend to have lower volatility than pretax returns have, but how much lower? This question is difficult to answer because after-tax returns are not generally available for analysis. Nevertheless, we can model the effect of taxes on risk and assess the subsequent impact on after-tax asset allocation. The basic idea is that for a given tax-inefficient asset class, the tax penalty is positively correlated with pretax return. Larger returns tend to trigger larger penalties, but the relationship is far from perfect.[21]

For simplicity, we assumed that the correlation between return and the tax penalty is constant across tax-inefficient asset classes. Using simulation, we examined the effect of this correlation under a range of assumptions.

Table 6. Implied Returns and Weights for Yale Endowment (High-Correlation Hedge Fund Index)

	Returns			Weights			
Asset Class	Pretax Implied Return	After-Tax Implied Return	Tax Haircut	Pretax	After-Tax Active Equity	After-Tax Indexed Equity	After-Tax Tax-Advantaged Equity
Absolute return	4.2%	3.3%	–0.9%	17.8%	0.0%	0.0%	0.0%
Active equity	9.2	7.1	–2.1	15.7	0.0	0.0	0.0
Indexed equity	9.2	8.7	–0.5	0.0	0.0	45.6	0.0
Tax-advantaged equity	9.2	10.6	1.4	0.0	0.0	0.0	25.0
Bonds	1.5	0.9	–0.7	4.9	35.0	25.8	25.6
Natural resources	10.2	9.2	–1.0	7.9	12.2	0.0	5.8
Real estate	11.4	9.9	–1.5	20.2	13.9	9.0	2.3
Private equity	11.0	9.8	–1.2	32.0	38.9	19.6	41.3
Cash	1.5	0.8	–0.7	1.5	0.0	0.0	0.0

Note: See notes to Table 3.

The details of our model are discussed in Appendix C, and the main results are shown here. **Table 7** shows the effect of positive correlation between pretax return and the tax penalty on tax-inefficient asset classes—absolute return (hedge funds), active equity, natural resources, real estate, and private equity.[22] As the correlation between return and the tax penalty increases, the volatilities of the tax-inefficient asset classes decline. The decline is not uniform across asset classes, however, on either an absolute basis or a percentage basis. For example, when $\rho = 0.75$, the volatility of absolute return (hedge funds) declines from 6.0% per year to 5.0% per year. The volatility declines by 1 percentage point (pp) in level, or by 16.7%. At the same time, the volatility of real estate declines from 23.9% per year to 21.5% per year.

Table 8 shows the impact of tax adjusting the covariance matrix on Yale's after-tax asset allocation when a tax-advantaged equity class is available. The far-left column in Table 8 is an exact replica of the far-right column in Table 3; it displays Yale's after-tax asset allocation when the pretax covariance matrix is used in the optimization. The columns to the right show how that asset allocation is modified as the correlation, ρ, between pretax return and the tax haircut increases to a perfect correlation ($\rho = 1$).

The result is higher allocations for tax-advantaged equity and real estate and lower allocations for absolute return (hedge funds), bonds, natural resources, and private equity. In summary, weight flowed from tax-inefficient asset classes whose volatilities were less diminished to tax-inefficient asset classes whose volatilities were more diminished. Overall, the net allocation to tax-inefficient asset classes increased. This change was facilitated by an increased allocation to tax-advantaged equity, which, again, offset part of the tax penalty.

Table 7. Change in After-Tax Asset Class Volatilities

		Change in Volatility Using After-Tax Covariance Matrix			
Asset Class	Volatility with Pretax Covariance Matrix	$\rho = 0.25$	$\rho = 0.50$	$\rho = 0.75$	$\rho = 1$
Absolute return	6.0%	–0.3	–0.7	–1.0	–1.3
Active equity	17.2	–1.0	–1.9	–2.8	–3.8
Indexed equity	17.2	0.0	0.0	0.0	0.0
Tax-advantaged equity	17.2	0.0	0.0	0.0	0.0
Bonds	4.6	0.0	0.0	0.0	0.0
Natural resources	23.5	–0.6	–1.1	–1.7	–2.2
Real estate	23.9	–0.8	–1.6	–2.4	–3.2
Private equity	20.6	–0.4	–0.8	–1.2	–1.5
Cash	0.1	0.0	0.0	0.0	0.0

Note: Change in volatility is measured in percentage points.

Table 8. After-Tax Asset Allocation for Yale Endowment (Low-Correlation Hedge Fund Index, After-Tax Covariance Matrix)

		Change in Weight Using After-Tax Covariance Matrix[a]			
Asset Class	Weight with Pretax Covariance Matrix	$\rho = 0.25$	$\rho = 0.50$	$\rho = 0.75$	$\rho = 1$
Absolute return	11.9%	–0.2	–0.6	–3.2	–6.2
Active equity	0.0	0.0	3.6	11.3	20.9
Indexed equity	0.0	0.0	0.0	0.0	0.0
Tax-advantaged equity	26.4	–0.2	1.0	3.9	7.6
Bonds	14.5	–2.5	–5.4	–8.4	–12.2
Natural resources	5.2	0.9	0.5	–1.3	–3.6
Real estate	2.8	3.0	6.1	9.7	13.8
Private equity	39.1	–1.3	–5.3	–12.2	–20.4
Cash	0.0	0.0	0.0	0.0	0.0

Note: Change in weight is measured in percentage points.

[a]Reflects the case presented earlier where tax-advantaged equity is available, shown in the far-right column in Table 3.

Estate or Charity Dispositions vs. Liquidation

In the examples we have considered, the tax drag for hedge funds makes them unattractive unless their correlation with equity is low and their tax impact can be offset. Private equity often gets a boost when taxes are introduced, owing to the assumptions about relative tax efficiency. If indexed equity is available, however, the optimizer lowers the exposure to private equity, but the optimizer raises this exposure again when tax-advantaged equity is available—as shown, for example, in Tables 3, 4, and 6. That pattern reflects the high correlation between private and public equity, but it is tempered by the tax haircut.

To this point, all our examples have assumed estate or charity dispositions at the end of the holding period. An alternative assumption is that assets will be liquidated at the end of the period. Under the assumption of liquidation, the tax penalty for the most inefficient asset classes would be lower, depending on the length of the holding period, and the benefit of a tax-advantaged equity strategy would also be lower. In modeling the benefit of a tax-advantaged equity strategy, the potential gain of 1.9% per year for an estate or charity disposition after 10 years would drop to 1.0% for the same 10-year period if all portfolios were assumed to be liquidated at the end of the period.

The likelihood of liquidation or estate disposition depends on each investor's specific goals for the use of the funds as well as the investor's total wealth. Although situations vary, it is more likely that an investor with a net worth of $1 million would need to spend down assets for current needs than an investor with a net worth of $1 billion, who may be keenly focused on estate planning, motivated by a desire to minimize the tax impact of intergenerational wealth transfer. Thus, the former might focus more on liquidation economics, whereas the latter might focus more on the economics of assets passing through an estate.

The Yale Model Adapted to the Needs of a Taxable Investor

In his second book, *Unconventional Success*, David Swensen (2005) explicitly recognized the harm of ignoring taxes once you change environments: "A serious fiduciary with responsibility for taxable assets recognizes that only extraordinary circumstances justify deviation from a simple strategy of selling losers and holding winners" (p. 217). The taxable world demands a radically different approach from the one taken in the tax-exempt world, and it requires tools like those described in our study.[23]

What are the lessons learned from evaluating the Yale endowment with these tools? First, it is essential to combine risk and tax considerations in the design stage of a portfolio. The tax effect cannot be layered in after the fact.

Second, in the after-tax version of the Yale endowment with only actively managed equity strategies, the optimizer includes no equity exposure at all. That might sound heretical to many managers, but our analysis indicates that if you are facing a tax haircut similar to the average one for active equity funds, then it is not worth owning any equities at all if you believe in Yale's asset allocation. Once you allow indexing, however, with its significantly reduced trading, then suddenly a good part of the portfolio looks like a simple 1960s-style 60%/40% stock/bond pension fund. Adding a tax-advantaged equity strategy goes one step further by allowing the portfolio to hold more tax-inefficient assets, such as hedge funds, taking advantage of their diversification benefits while offsetting some of the tax drag.

Third, correlations matter even more in the presence of taxes. In our analysis, hedge funds that are highly correlated with equity markets never are included in the Yale portfolio once taxes are introduced. This finding is an indication that there is no place for highly inefficient tax strategies unless there is a big diversification benefit. Alternative strategies that are highly correlated with equities—and that generate a lot of gains from trading—cannot survive in a portfolio designed to maximize after-tax returns.

Fourth, the introduction of an after-tax covariance matrix, with lower volatilities for tax-inefficient asset classes, strengthens the results. As tax-inefficient asset classes become more attractive, a greater allocation to tax-advantaged equity is required to offset the tax drag.

The new framework and the accompanying metrics can help investors and advisers pursue many of the benefits of the endowment model in a way that takes into account the harsh after-tax reality of the world of UHNW investors. So what would Yale do if it were taxable? A world-class shop like Yale would integrate tax considerations into its asset allocation process from the beginning rather than layering them in after the fact.

Appendix A. Indexes Used as Asset Class Proxies

In **Exhibit A1**, we show the indexes that we used as asset class proxies.

Exhibit A1. List of Asset Classes and Benchmarks

Proxy No.	Asset Class	Benchmark
1	Low-correlation absolute return	Credit Suisse AllHedge Index
2	High-correlation absolute return	HFRI Fund Weighted Composite Index
3	World equity, active	MSCI ACWI Index
4	World equity, tax neutral	MSCI ACWI Index
5	World equity, tax advantaged	MSCI ACWI Index
6	Fixed income	Barclays U.S. Aggregate Bond Index
7	Municipal bonds	Barclays U.S. Municipal Bond Index
8	Natural resources	S&P North American Natural Resources Index
9	Private equity	Russell 2000 Index
10	Real estate	Dow Jones U.S. Real Estate Index
11	Cash	BofA Merrill Lynch 0–3 Month U.S. Treasury Bill Index

Note: For descriptions of these indexes, see Appendix D.

Appendix B. Analysis of NACUBO Portfolio

When the starting point is the pretax allocation for the average endowment from NACUBO, our analysis yields after-tax weights fairly similar to those for the Yale endowment, as shown in **Table B1**. The high tax penalty for active management prevents the NACUBO model from holding any equity at all when only active equity is available because private equity provides much of the same exposure, owing to the assumption that it has considerably lower turnover. The After-Tax Indexed Equity column shows how equity provided through traditional indexing generates an outcome that looks very similar to the old pension standby of 60% stocks and 40% bonds, with a light sprinkling of alternatives. The After-Tax Tax-Advantaged Equity column shows the results when constrained tax-advantaged equity is an available strategy, which allows for more alternatives because of the presence of tax-advantaged equity. In addition, it reflects the odd outcome of actively managed equity selected over traditional indexing. That counterintuitive result reflects the fact that the estimated tax benefit from tax-advantaged equity happens to be larger than the penalty for active management, meaning that a taxable investor would be better off holding more active equity because it allows more tax-advantaged equity.

Table B1. Implied Returns and Weights for NACUBO (Low-Correlation Hedge Fund Index)

	Returns			Weights			
Asset Class	Pretax Implied Return	After-Tax Implied Return	Tax Haircut	Pretax	After-Tax Active Equity	After-Tax Indexed Equity	After-Tax Tax-Advantaged Equity
Absolute return	2.1%	1.6%	–0.5%	12.0%	0.0%	0.0%	11.0%
Active equity	9.3	7.1	–2.1	41.0	0.0	0.0	9.7
Indexed equity	9.3	8.8	–0.5	0.0	0.0	65.7	0.0
Tax-advantaged equity	9.3	10.7	1.4	0.0	0.0	0.0	26.7
Bonds	1.6	0.9	–0.7	15.0	44.0	30.9	21.9
Natural resources	10.5	9.5	–1.0	10.0	20.5	2.0	11.8
Real estate	10.5	9.1	–1.4	10.0	8.1	1.0	0.0
Private equity	10.2	9.1	–1.1	10.0	27.5	0.4	19.0
Cash	1.5	0.8	–0.7	2.0	0.0	0.0	0.0

Notes: See notes to Table 3. Alternatives are assumed to be weighted equally among natural resources, real estate, and private equity because NACUBO does not provide that breakdown.

Sources: Pretax weights are from www.nacubo.org; returns and after-tax weights were calculated by Aperio Group.

Appendix C. Adjusting a Covariance Matrix for Taxes

We simulated the history of after-tax returns by assuming that the pair (r, h), where r represents pretax returns and h represents tax haircuts, is joint normally distributed and the after-tax return in period t is given by

$$R_t = r_t - h_t.$$

In order to simulate after-tax returns using the history of pretax returns, we generated a sequence of tax haircuts by drawing from the conditional distribution of h given $r = r_t$:

$$h_t \,|r_t \sim N\left[\mu_h + \rho\frac{\sigma_h}{\sigma_r}\left(r_t - \mu_r\right), \left(1-\rho^2\right)\sigma_h^2\right].$$

Given an estimate of the volatility of r, σ_r, we estimated the volatility of h:

$$\sigma_h = \sigma_r \frac{\mu_h}{\mu_r},$$

where μ_h is the expected tax penalty for each asset class, calculated as outlined in the article, and μ_r is the expected pretax return for each asset class. Given the assumption for σ_h, the conditional distribution of h given $r = r_t$ becomes

$$h_t \,|r_t \sim N\left[\mu_h + \rho\frac{\mu_h}{\mu_r}\left(r_t - \mu_r\right), \left(1-\rho^2\right)\sigma_r^2\frac{\mu_h^2}{\mu_r^2}\right].$$

For each asset class and each period, we took 1,000 draws from the conditional distribution, calculated an after-tax return for each draw, and then averaged over the 1,000 simulated after-tax returns. The resulting tax-adjusted return series for all asset classes was then used to estimate an after-tax covariance matrix.

Tables 7 and 8 display the results for various values of the correlation parameter, ρ, when we replaced the pretax covariance matrix with the after-tax covariance matrix (estimated using that value of ρ) in the final optimization step.

Appendix D. Index Descriptions

In **Exhibit D1**, we provide descriptions of the indexes we used as benchmarks.

Exhibit D1. Descriptions of Indexes Used as Benchmarks

Benchmark	Description
Credit Suisse AllHedge Index	The Credit Suisse AllHedge Indexes are designed to provide transparent, representative, and objective benchmarks of the 10 style-based investment strategies of the hedge fund universe. The unique features of the indexes enable them to serve as the underlying benchmark for a broad suite of investment products. The composite index of the family is the Credit Suisse AllHedge Index. The subindexes covering the individual investment strategies are known collectively as the Credit Suisse AllHedge Strategy Indexes.
HFRI Fund Weighted Composite Index	• Includes over 2,200 constituent funds • Includes both domestic and offshore funds • Equal-weighted index • All funds report assets in US dollars • No funds of funds included in the index • All funds report net-of-all-fees returns on a monthly basis • Funds have at least $50 million under management or have been actively trading for at least 12 months
MSCI ACWI Index	This index is a free-float-adjusted market capitalization–weighted index that is designed to measure the equity market performance of developed and emerging markets. The MSCI ACWI consists of 44 country indexes: 23 developed and 21 emerging market country indexes. The developed market country indexes are Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany, Hong Kong, Ireland, Israel, Italy, Japan, the Netherlands, New Zealand, Norway, Portugal, Singapore, Spain, Sweden, Switzerland, the United Kingdom, and the United States. The emerging market country indexes are Brazil, Chile, China, Colombia, Czech Republic, Egypt, Greece, Hungary, India, Indonesia, Malaysia, Mexico, Peru, the Philippines, Poland, Russia, South Africa, South Korea, Taiwan, Thailand, and Turkey. Net of foreign withholding taxes; dividends are reinvested.
Barclays U.S. Aggregate Bond Index	This index is a broad-based flagship benchmark that measures the investment-grade, US dollar–denominated, fixed-rate taxable bond market on a total-return basis. The index includes Treasuries, government-related and corporate securities, mortgage-backed securities (agency fixed-rate and hybrid adjustable-rate mortgage pass-throughs), asset-backed securities, and commercial mortgage-backed securities (agency and non-agency).
Barclays U.S. Municipal Bond Index	This index covers the US dollar–denominated, long-term tax-exempt bond market. The index has four main sectors: state and local general obligation bonds, revenue bonds, insured bonds, and pre-refunded bonds.

(continued)

Exhibit D1. Descriptions of Indexes Used as Benchmarks (continued)

Benchmark	Description
S&P North American Natural Resources Index	This index is a total-return-weighted index of 146 stocks that represent US traded securities that are classified under the GICS[a] energy and materials sector, excluding the chemicals industry and the steel subindustry.
Russell 2000 Index	This index measures the performance of the small-cap segment of the US equity universe on a total-return basis. The Russell 2000 Index is a subset of the Russell 3000 Index that represents approximately 10% of the total market capitalization of that index. It includes approximately 2,000 of the smallest securities based on a combination of their market capitalizations and current index membership. The Russell 2000 is constructed to provide a comprehensive and unbiased small-cap barometer and is completely reconstituted annually to ensure that larger stocks do not distort the performance and characteristics of the true small-cap opportunity set.
Dow Jones U.S. Real Estate Index	The objective of this index is to represent REITs and other companies that invest directly or indirectly in real estate through development, management, or ownership, including property agencies. The index is a subset of the Dow Jones U.S. Market Index, which covers 95% of US securities based on float-adjusted market capitalization.
BofA Merrill Lynch 0–3 Month U.S. Treasury Bill Index	This index tracks the performance of US dollar–denominated US Treasury bills publicly issued in the US domestic market with a remaining term to final maturity of less than three months.

[a]GICS is a registered trademark of MSCI and Standard & Poor's.

Sources: The edited definitions in this exhibit are from the Morningstar Principia database except for those of the S&P North American Natural Resources Index (S&P 500 Index factsheet) and the BofA Merrill Lynch 0–3 Month U.S. Treasury Bill Index (SEC data).

Notes

1. There are many dimensions to tax-aware investing, and Wilcox, Horvitz, and diBartolomeo (2006) covered a broad range of issues related to after-tax investing. They considered the role of mean–variance optimization in after-tax asset allocation, the difficulties associated with concentrated positions, and the location problem, which concerns which assets should be held in tax-deferred accounts. The location problem is addressed in Dammon, Spatt, and Zhang (2004), who found an inverse relationship between equity allocation and the fraction of assets held in a tax-deferred account. Berkin and Ye (2003) used Monte Carlo simulation to assess the impact of HIFO (highest in, first out) asset management and loss harvesting. Reichenstein, Horan, and Jennings (2012) view taxable accounts as partnerships between an investor and the government, and this perspective has implications for optimal asset location and the decision to allocate to a traditional or a Roth IRA.
2. The success of the Yale and other endowments has attracted interest throughout the finance community, including advisers to UHNW investors. Some investment advisers' websites explicitly recommend the Yale endowment model to UHNW clients. For example, Crestone Capital Advisors, LLC, boasts on its website (www.crestonecap.com), "Crestone employs an investing strategy similar to endowments like Yale, Harvard, and Stanford. This has helped our clients not only protect and grow their wealth, but also reduce the impact of the recent public equity and bond market volatility." Mladina and Coyle (2010) asked whether the success of the Yale model comes from asset exposures or manager skill and discussed the Yale model from the perspective of a high-net-worth investor, but there is no mention of taxes in their article. The Yale model has also been emulated by pensions, as noted by Morgenson (2014).
3. In a study of endowment performance, Barber and Wang (2013) attributed the superior returns of the Yale and other Ivy League endowments to asset allocation rather than manager selection. Specifically, they did not find statistically significant alpha in a regression of endowment returns on traditional and alternative asset class returns. This finding is apparently at odds with Swensen's (2005) emphasis on the importance of manager selection. At least part of the discrepancy may be explained by the indexes used to represent alternative asset classes. For example, Barber and Wang used returns to the HFRI Fund Weighted Composite Index as a proxy for hedge fund returns. This index is notorious for its survivorship bias, so it is arguable that the impressive index returns are largely due to the most successful managers. Barber and Wang attempted to address the issue with alternative hedge fund proxies. However, the opacity and incompleteness of hedge fund data make it difficult to draw reliable conclusions.
4. Citing the empirical literature, Swensen (2005) reported that turnover in a typical actively managed mutual fund diminishes pretax return by roughly 2% per year. Estimates of an approximately 2% tax haircut for active equity can be found—for example, Jeffrey and Arnott (1993); Swensen (2005, p. 258); Geddes (2011). Poterba (2001, p. 92) suggested that taxable investors in stocks might lose as much as 3.5% per year.
5. There are legal restrictions governing the structure of tax-efficient equity. One vehicle is a separately managed account, or SMA. More information is available in Wilcox et al. (2006) and Geddes (2011).
6. Persson and Mustillo (2012) estimated the impact of taxes on Yale's expected returns.
7. The advisory world potentially faces a subtle conflict of interest in wanting to be perceived as adding value with investment products that may not look as good after taxes as they do before taxes. Raising client awareness about the tax penalty can prove challenging because of both the complexity of the subject and the pall it can cast on an otherwise exciting story for some of the most glamorous strategies. This agency problem reflects the asymmetry of investors paying the cost of adviser-designed portfolios that do not fully incorporate taxes. As Dan diBartolomeo (2010) put it, "Although it is clear to everyone in the industry that after-tax returns are what economically matter to taxable investors, many firms continue to take the view that 'we can ignore taxes entirely as long as our competitor managers do the same.' In our view, this position of intentionally providing less than the best available services

to separate account clients is unsustainable in the long run and borders on a breach of fiduciary responsibility" (p. 6).

8. See, for example, Merton (1980).
9. As noted by El Karoui (2013), "There seems to be a consensus in the finance literature that the estimation of the covariance is 'easy' and that the more difficult aspect of the Markowitz problem (and other portfolio optimization problems) is hard. By contrast, statisticians working in high-dimensional inference have recently devoted a lot of effort to improving covariance estimation, which is thought to be a hard task" (pp. 742–743).
10. Yale University (2013).
11. Data are as of 30 June 2013.
12. Our covariance matrix calculation is based on monthly data weighted exponentially with a 60-month half-life. The sample period is January 1999–June 2013.
13. In some situations, it is possible to avoid income tax entirely (e.g., with assets passed through an estate or donated to charities).
14. For details on the value added from a tax-advantaged strategy, see Geddes (2011).The potential benefit from a tax-advantaged strategy reflects the highest applicable federal income tax rates as of 31 December 2013.
15. The harvesting opportunities do tend to diminish significantly after 10 years, but the deferral can continue to add value even after much of the harvesting has already been realized.
16. Taxable investors in California, for example, can benefit by as much as 2.40% per year for portfolios passed through an estate and 1.20% for those liquidated at the end of a 10-year period. The marginal income tax rate in California is assumed to be 13.3% and fully deductible against federal income.
17. Investors can also achieve commodity exposure through natural resource stocks, which can be managed with a high degree of tax efficiency.
18. The assumption on the breakdown between long- and short-term gains was kept consistent across all asset classes except for private equity. For empirical results based on older data, see Bogle (1999). Bogle's estimate from October 1999 was 65% long-term gains and 35% short-term gains.
19. All tax adjustments are based on the highest tax bracket rates for US federal income tax only. Those rates reflect tax law changes effective 1 January 2013, which include a top short-term gains rate of 39.6%, a long-term gains rate of 20.0%, an additional 3.8% tax on net investment income (the Medicare tax), and an effective additional 1.2% rate increase due to the phaseout of deductions for top-bracket taxpayers (the Pease phaseout). All these rates combine to 44.6% for short-term gains and ordinary income and 25.0% for long-term gains and dividend income. All dividend income has been assumed to come from qualified dividends. State taxes would make the tax haircut even more expensive, but the outcomes remain similar to what was calculated for the impact of only federal income tax.
20. For the highest bracket in California, the marginal tax rate includes the 44.6% federal rate (as calculated in the previous endnote) plus the top California rate of 13.3%. The California liability is assumed to be deductible for federal tax purposes (conventional 39.6% plus 3.8% tax on net investment income, though the latter is subject to some limitations on deductibility of state taxes). The gross California rate of 13.3% thus reduces to a net impact of $(1 - 0.396 - 0.038) \times 0.133 = 7.5\%$, which means a total liability of 52.1% when combined with the federal rate of 44.6%.
21. The correlation between the tax penalty and the return can depend on the legal structure of how an economic exposure is delivered. As mentioned previously, for natural resources and other such asset classes, some strategies, such as those using futures contracts, can incur a significant tax drag whereas others, such as direct ownership of timber or extractive industries, may take advantage of tax code features that favor such investment structures. The more the gains are distributed, the higher the correlation. Some structures, such as certain hedge funds, have at times delivered the emotionally unpleasant combination of pretax losses and tax liability. In other words, investors sometimes face the outrage of negative pretax returns made even worse by having to pay tax on distributed gains for the same period during which they are losing money. Sophisticated research in behavioral finance is not required to understand that such phenomena tend to irritate investors.
22. We did not adjust the volatility of indexed equity, bonds, or cash because any tax adjustment to these asset classes is consistent over time. We did not adjust the volatility of tax-advantaged equity because the scant evidence we have seen indicates no correlation between pretax return and tax alpha.
23. Mean–variance optimizers can show how risk, return, and taxes all combine within a portfolio. Measuring how taxes change the asset allocation provides valuable insight into just how expensive the tax penalty can be for investments that trigger significant tax liability. However, investors should avoid viewing output from a mean–variance optimizer as a definitive best solution. Investment decisions require judgment of many factors around return, risk, and taxes, and the output from an optimizer should not be viewed as "the right answer." Furthermore, the assumptions that went into the particular analyses described in this article and the taxonomy of categorizing described herein may not exactly match a specific investor situation. The taxonomy of categorizing different investment strategies into asset classes helps streamline decision making, but it can lead to oversimplification. For example, the large number and different types of hedge funds mean that a particular representative index may not serve as a satisfactory proxy for every hedge fund available to a UHNW investor. Similarly, although hedge funds are generally perceived as relatively inefficient in terms of tax management, some funds may be quite efficient, and others do distribute significant taxable income, often at the high tax rates applied to short-term gains. Nonetheless, neither the complexity nor the lack of precise data should be construed as excuses to avoid serious analysis of how returns, risk, and taxes combine within a portfolio.

References

Barber, Brad M., and Guojun Wang. 2013. "Do (Some) University Endowments Earn Alpha?" *Financial Analysts Journal*, vol. 69, no. 5 (September/October): 26–44.

Berkin, Andrew L., and Jia Ye. 2003. "Tax Management, Loss Harvesting, and HIFO Accounting." *Financial Analysts Journal*, vol. 59, no. 4 (July/August): 91–102.

Black, Fischer, and Robert Litterman. 1992. "Global Portfolio Optimization." *Financial Analysts Journal*, vol. 48, no. 5 (September/October): 28–43.

Bogle, John C. 1999. "The Economics of the Mutual Fund Industry: For Fund Investors … for Fund Managers." Presentation at Albright College, Reading, PA (5 October).

Dammon, Robert M., Chester S. Spatt, and Harold H. Zhang. 2004. "Optimal Asset Location and Allocation with Taxable and Tax-Deferred Investing." *Journal of Finance*, vol. 59, no. 3 (June): 999–1037.

diBartolomeo, Dan. 2010. "Northfield for Everyone: Analytics for Private Wealth." *Northfield News* (September): 1–7 (www.northinfo.com/documents/390.pdf).

El Karoui, Noureddine. 2013. "On the Realized Risk of High-Dimensional Markowitz Portfolios." *SIAM Journal on Financial Mathematics*, vol. 4, no. 1: 737–783.

Geddes, Patrick. 2011. "The Pursuit of After-Tax Returns: Indexed Exchange Traded Funds versus Indexed Separately Managed Accounts." *Investments & Wealth Monitor* (September/October): 38–43.

Jeffrey, Robert H., and Robert D. Arnott. 1993. "Is Your Alpha Big Enough to Cover Its Taxes?" *Journal of Portfolio Management*, vol. 19, no. 3 (Spring): 15–25.

Merton, Robert C. 1980. "On Estimating the Expected Return on the Market: An Exploratory Investigation." *Journal of Financial Economics*, vol. 8, no. 4 (December): 323–361.

Mladina, Peter, and Jeffery Coyle. 2010. "Yale's Endowment Returns: Manager Skill or Risk Exposure?" *Journal of Wealth Management*, vol. 13, no. 1 (Summer): 43–50.

Morgenson, Gretchen. 2014. "Slamming a Door on Hedge Funds." *New York Times* (20 September): www.nytimes.com/2014/09/21/business/slamming-a-door-on-hedge-funds.html?_r=0.

Persson, Jon, and Mike Mustillo. 2012. "The Endowment Model and the Philanthropic Taxable Investor." White paper, Geller Family Office Services LLC.

Poterba, James M. 2001. "Taxation, Risk-Taking, and Household Portfolio Behavior." NBER Working Paper 8340 (June).

Reichenstein, William, Stephen M. Horan, and William W. Jennings. 2012. "Two Key Concepts for Wealth Management and Beyond." *Financial Analysts Journal*, vol. 68, no. 1 (January/February): 14–22.

Sharpe, William F. 1974. "Imputing Expected Security Returns from Portfolio Composition." *Journal of Financial and Quantitative Analysis*, vol. 9, no. 3 (June): 463–472.

Swensen, David F. 2005. *Unconventional Success: A Fundamental Approach to Personal Investment*. New York: Free Press.

Wilcox, Jarrod, Jeffrey E. Horvitz, and Dan diBartolomeo. 2006. *Investment Management for Taxable Private Investors*. Charlottesville, VA: CFA Institute Research Foundation.

Yale University. 2013. "The Yale Endowment: 2013." Annual report of the Yale endowment (http://investments.yale.edu/images/documents/Yale_Endowment_13.pdf).