Principles of after-tax asset allocation

by Stephen M. Horan, CFA Institute

The notion of after-tax asset allocation is gaining acceptance among private wealth managers. This article presents practical methods of calculating an investor's after-tax asset allocation, particularly as it relates to taxable accounts. The after-tax value of a taxable account can be substantially less than its stated pretax value, especially for long time horizons. As a result, a client's pretax asset allocation can be a poor proxy for her more economically meaningful after-tax asset allocation, especially when assets are held in accounts with different tax treatments. The notion of after-tax asset allocation in turn has implications for after-tax portfolio optimisation and asset location strategies.

Introduction

Private wealth managers have traditionally used pretax values to calculate an investor's asset allocation, making no distinction between assets held in tax-sheltered accounts and taxable accounts. However, the same asset held in different types of accounts can produce different after-tax cash flows, have different after-tax risk profiles, and hence have different after-tax values.

Pension fund managers can often abstract from the tax implications of their investment decisions. High net worth (HNW) individuals, on the other hand, are unable to consume pretax wealth and must therefore make consumption decisions in an after-tax environment. As a result, private wealth managers study tax-efficient investing techniques. Much of their attention has focused on capital gain recognition, wealth transfer strategies and estate planning issues, which have important and often obvious tax implications. Less attention has focused on equally important but more latent issues, such as how taxes affect risk and hence an investor's effective after-tax asset allocation. These issues in turn have implications for tax-efficient portfolio optimisation and asset location (i.e. the practice of placing certain types of assets in certain types of accounts).

Tax-efficient investing is a global concern for HNW investors for several reasons. First, an investment's tax drag depends on its asset class and how it is traded. For example, all but two of 24 industrialised and developing countries studied by the American Council for Capital Formation (ACCF) in 1998 have different tax rates for interest income and capital gains, including Canada, Germany, Italy, France, the Netherlands, Sweden, the UK, and Japan. Five of the countries surveyed have holding period requirements to qualify for lower capital gains tax rates. As a result, tax treatment can depend on trading behaviour and portfolio turnover.

Moreover, most industrialised and developing countries have tax incentives to encourage retirement savings. A related ACCF study indicates that two-thirds of the countries surveyed offer tax-advantaged savings accounts. Examples include Australia, Canada, Germany, Italy, the Netherlands, and the UK. And all but one of these countries permits tax-deductible contributions. In addition, defined-contribution retirement plans (whether sponsored by an employer or a state) and deferred employee compensation plans have a similar global presence. Some 30% of the countries in the ACCF study, for example, have some form of state-sponsored defined-contribution retirement plan.

Substantial sums of money have accumulated over time in various tax-sheltered investment vehicles. Broadly speaking, the capitalised value of pension benefits in defined-benefit plans is a tax-deferred asset for pensioners. Considered in this general context, an understanding of tax-efficient portfolio management should consider how different asset classes and tax-sheltered accounts are taxed.

Models developed in a pretax environment do not necessarily apply in a more economically relevant after-tax environment, and this tenet certainly applies to asset allocation. Advisors relying on pretax values to determine a client's mix of equity, fixed income, and other assets can be misled about a portfolio's true risk profile. Depending on an investment's asset class and the type of account in which it is held, the taxing authority may share in a portion of an investor's return and hence absorbs some of her investment risk. Alternatively, the taxing authority may effectively only participate in the

investment's principal value. Regardless, either circumstance affects after-tax asset allocation.

This article introduces a fundamental approach to converting pretax values to after-tax values. By considering the basic forms of taxation present in most developed and developing countries, it captures the heuristic implications of various types of investment vehicles in various tax jurisdictions.

Different types of accounts

Consider three types of investment accounts. The first is tax-deferred accounts, or TDAs. Contributions to these accounts are often made on a pretax basis, and the investment returns accumulate on a tax-deferred basis until funds are withdrawn at which time they are taxed as ordinary income. As such, they are sometimes said to have front-end loaded tax benefits. Examples include traditional IRAs in the US or optional Tier II accounts in the UK's privatised Social Security system, but there are many others.

A second class of accounts has back-end loaded tax benefits. These accounts are tax-exempt on prospective basis because, although contributions are taxable, their earnings can accumulate tax free even as funds are withdrawn. Examples include Roth IRAs in the US.

A third type of account is taxable accounts for which deposits are made on an after-tax basis. Investment returns in taxable accounts are taxed depending on the type of investment and the trading activity. However, at least a portion of return is usually taxed on an annual basis.

Therefore, the same asset is taxed differently depending on the type of account in which it is held. Because of their different tax treatment, balances held in these three types of accounts are not economically comparable, and calculating an investor's asset allocation based on pretax values can be quite misleading.

After-tax values

The concept of after-tax values can be easily understood using fundamental time value of money concepts. First, consider an investor who owns £500,000 of bonds held in a TDA. If the bonds earn 7% per annum, the account value will grow to over £701,000 in five years. However, if the funds are withdrawn at that time and taxed at a 40% tax rate, the after-tax cash flow is only about £420,000 in five years. To compute the present value of this after-tax cash flow, we need to know the correct discount rate. Because investment returns in a TDA are tax-deferred, the investor bears all the risk associated with the investment return, and the correct discount rate is the required return on the bonds. Assuming the realised return on the bonds is equal to the required return, the correct discount rate is 7%, and the present value is £300,000. This figure represents the after-tax value of the bonds held in the TDA, and its derivation is presented in the first row of Figure 1.

It is noteworthy, yet not coincidental, that the

Figure 1: An example of after-tax asset allocation

Account and security	Current pretax value	%	Future after-tax cash flow	Current after-tax value	%
TDA Bonds	500,000	33%	420,766	300,000	24%
Roth Equity	500,000	33%	805,255	500,000	40%
Taxable Equity	500,000	33%	669,113	463,915	37%
Total	**1,500,000**			**1,263,915**	

Figure 2: After-tax value factors for various types of investment accounts

Account type	After-tax ending wealth	Discount rate	After-tax present value of a pretax dollar
TDA	$(1 + r)^n (1 - t)$	r	$\frac{(1 + r)^n (1 - t)}{(1 + r)^n} = (1 - t)$
Roth IRA-type	$(1 + r)^n$	r	$\frac{(1 + r)^n}{(1 + r)^n} = 1$
Taxable Account	$[1 + r (1 - t)]^n$	$r_f + (1 - t)\ \beta_i\ [MRP]$	$\frac{[1 + r (1 - t)]^n}{(1 + r_f + (1 - t)\ \beta_i\ [MRP])^n} < 1$

£300,000 after tax value happens to equal £500,000 times (1 – 40%). For TDAs, the after-tax value is always the pretax value times 1 minus the tax rate regardless of the investment's return, the investment horizon, or the withdrawal pattern. In other words, the after-tax value of a TDA is equal to its pretax value times (1 – *t*) where *t* is the tax rate applied to withdrawals. The first row of Figure 2 summarises this result more generally for a dollar in a TDA.

We can apply similar logic to investment accounts with tax-free withdrawals, like the Roth IRA. Assume our investor also holds £500,000 of equity yielding 10% return per annum in this type of tax-sheltered account. The after-tax cash flow in five years is about £805,000, which also equals the pretax cash flow because withdrawals are not taxed (see the second row of Figure 1). Like the TDA, investors bear all investment risk because returns are not taxed either. Therefore, this future cash flow should also be discounted at 10%, which of course is the same rate we used to compound our initial balance. So, the after-tax value of tax-sheltered accounts similar to the Roth IRA is simply equal to its pretax value because neither returns nor withdrawals are taxed. Balances in these accounts are, by definition, after-tax balances.

The second row of Table 2 presents this result more generally. It illustrates that the after-tax values of a dollar in a TDA and a Roth IRA differ by a factor of (1 – *t*). Consequently, assets held in TDAs are always less valuable than the equivalent asset in a Roth-type account and should therefore be weighted less heavily when calculating an investor's asset allocation.

Dealing with taxable accounts is a bit more nuanced. Assets held in taxable accounts can be taxed in a variety of ways, depending on the asset class (e.g. stocks versus bonds), the type of trading (e.g. long-term gains versus short-term gains), and taxing jurisdiction. For simplicity, however, let us consider that our investor has another £500,000 of actively-managed equity in a taxable account with 100% turnover so that investment gains are short-term in nature and taxed at a 40% marginal tax rate – the top marginal tax rate in the UK. If the Capital Asset Pricing Model (CAPM) correctly models equity returns and the risk-free rate (r_f) is 4%, the portfolio beta (β) equals one, and the market risk premium (MRP) is 6%, then the pretax expected return is 10%. The after-tax return is 6% per annum; that is, 10% times (1 – 40%). Compounding our £500,000 for five years at 6%, the taxable account grows to about £669,000.

What is the rate by which this future after-tax cash flow should be discounted? Unlike the tax-sheltered accounts discussed above, the government shares investment risk with the investor because investment returns are taxed annually. This form of taxation is nearly identical to many private profit-sharing arrangements. As a result, the investor does not shoulder all of the risk associated with the investment – small solace for having to give up a portion of the investment returns but important for determining after-tax asset allocation.

One might understandably think that the correct discount rate is 6%, or 10% times (1 – 40%), the same as the after-tax return. But this intuition overlooks the fact that the equity portfolio's return has two components – a riskless return of 4% plus a risk premium of 6%. Although government taxes both components of the investments' return, its risk-sharing role relates only to the risk premium by definition. Therefore, the tax adjustment for the discount rate applies only to the risk premium. In this case, the risk-free rate is 4% and the after-tax risk premium is 6% times (1 – 40%), or 3.6%. The discount rate is the sum of these two figures, or 7.6%, yielding a present value of about £464,000 (see the third row of Figure 1).

The third row of Figure 2 presents this approach more generally. It demonstrates that the after-tax value of a taxable account is always less than the pretax value. The tax drag on the rate at which principal compounds in the numerator is heavier than the tax drag on the discount rate in the denominator. Therefore, the after-tax discount rate is always greater than the after-tax rate of compounding, and the after-tax present value factor is always less than one.

It can be proven mathematically that only the risk premium in the discount rate should have a tax adjustment, but an example will illustrate more clearly why this is so. Suppose both the taxable account and the Roth IRA hold risk-free securities yielding 4% rather than an equity portfolio. At a 40% tax rate, the respective after-tax values of the two accounts five years from now are about £563,000 and £608,000. Both investments are riskless, so their future after-tax cash flows should be discounted at the same risk-free rate, namely 4%. At this discount rate, the after-tax present values are about £463,000 and £500,000, respectively.

The after-tax value of the Roth IRA is of course equal to its pretax value. The taxable account has a smaller after-tax value, however, because it produces less cash flow for the same amount of risk. If the taxable account's cash flows were discounted by applying the tax adjustment to the risk-free rate, its after-tax value would equal that of the Roth IRA, which cannot be. The implication is that the tax

adjustment in the discount rate applies only to the risk premium.

In most countries, equity returns are taxed in a more complex manner than described in this illustrative example. The tax code in most countries, for example, distinguishes between dividends, short-term capital gains, and long-term capital gains. Models exist to accommodate these more realistic taxation schemes as well as non-CAPM measures of expected return. But the following principles remained unchanged:

1. The after-tax value of a TDA equals $(1 - t)$ times its pretax value.
2. The after-tax value of a Roth-type account equals its pretax value.
3. The after-tax value of a taxable account is always less than its pretax value, and the difference grows with time.

Let us revisit our investor who has a total of £1.5m evenly distributed across the three accounts on a pretax basis (see Figure 1). Using a traditional view of asset allocation, one would conclude that two-thirds of her assets are invested in equity and one-third is invested in fixed income. An after-tax approach that recognises these three principles suggests instead that less than a quarter of her assets are invested in fixed income. In this case, a pretax approach substantially underestimates the risk of the portfolio.

Portfolio optimisation and asset location

We have seen how basic principles of measuring asset allocation in a pretax environment do not necessarily apply in a more economically relevant after-tax environment. The same holds true for portfolio optimisation techniques. In other words, pretax efficient frontiers are not reasonable proxies for after-tax efficient frontiers. Portfolio optimisation in an after-tax framework requires a series of adjustments. Merely substituting after-tax returns for pretax returns in the optimisation algorithm, however, is insufficient because investments often have large taxable events at the end of an investment horizon when they are liquidated or withdrawn from an account.

In a pretax framework, a portfolio manager might implement an algorithm to find the pretax asset allocation that maximises portfolio return subject to a risk constraint (often in the form of a given portfolio standard deviation). In an after-tax framework, the portfolio manager would first recognise that after-tax values, rather than pretax values, produce the economically relevant asset allocation. These after-tax values incorporate taxable events triggered by liquidating an asset or withdrawing funds from a TDA at the end of an investment horizon. The after-tax values resulting from an optimisation process would be converted to pretax values to implement the proposed asset allocation.

Second, an asset's annual after-tax return would be substituted for its pretax return. This step is straightforward using the example above, but grows more complex when returns in the taxable account are subject to multiple forms of taxation, such as ordinary income, realised short-term capital gains, realised long-term capital gains, and deferred capital gains. Researchers have nonetheless developed measures of effective after-tax return that incorporate these factors, so this adjustment can be implemented in more sophisticated taxing regimes.

Third, the portfolio manager would substitute the asset's after-tax standard deviation of returns for pretax standard deviations in the optimisation algorithm. This adjustment is based on the notion that government shares in the investment risk of assets held in taxable accounts. The after-tax standard deviation for both the TDA and the Roth-type account is simply equal to the pretax standard deviation because returns are not taxed and the investor bears all the risk of her investments. For taxable accounts, however, it can be demonstrated that the after-tax standard deviation is equal to the pretax standard deviation times $(1 - t)$. This relationship changes for more complex taxing regimes but is conceptually similar. With these modifications, private wealth managers can more efficiently optimise the disposition of their clients' assets.

Because after-tax values, after-tax returns, and after-tax standard deviations depend on the type of account in which the asset is held, the same asset held in different accounts is essentially a distinct after-tax asset. In other words, an investor optimising between two different asset classes (e.g. stocks and bonds) across three types of accounts has six different after-tax assets to allocate – stocks or bonds in each of the three accounts. The optimisation process must be constrained by the pretax balances available in each of the accounts.

Such an after-tax portfolio optimisation process tends to locate stock in taxable accounts and bonds in tax-advantaged accounts (as we have in our example) because stock is generally more tax efficient than fixed income. Treating a specific asset-account combination as a unique investment in a portfolio optimisation process will locate assets across different accounts in a tax efficient manner. As a result, after-tax portfolio optimisation and asset location are synonymous concepts.

Conclusion

Despite a long history of taxing investment returns, most modern portfolio theory is grounded in a pretax framework. The profession has yet to embrace specific models of after-tax asset allocation and portfolio optimisation for several reasons. First, incorporating the impact of taxes in an intuitive and applicable way has to date been illusive. Second, the cost of ignoring taxes in portfolio management (at least as it relates to risk) is often times latent. For example, the outcome of asset allocation and asset location strategies not chosen by an advisor are unobservable or at least ambiguous to investors and their advisors. As a result, a client rarely sees the after-tax outcome of a more tax-efficient disposition of assets. Third, pension funds, the profession's largest clients are tax exempt so researchers and portfolio managers could justifiably ignore taxes. The growing importance of private wealth in the investment management industry increases the impetus to incorporate tax considerations into asset allocation decisions.

The academy is still grappling with developing and promoting tractable, intuitive after-tax portfolio management models that financial advisors can readily use in their daily work. However, private wealth advisors can make reasonably accurate adjustments to traditional measures of asset allocation using relatively simple techniques.

Biography:

Stephen Horan, Ph.D., CFA is Head of Private Wealth in the Education Division of CFA Institute. Prior to joining CFA Institute, Dr. Horan was Professor of Finance at St. Bonaventure University, a Financial Analyst and Forensic Economist in private practice, a Principal of Alesco Advisors LLC, an Account Manager with Manning & Napier Advisors, Inc., and a Stockbroker with Quick and Reilly, Inc. in New York, NY.

Dr. Horan has co-authored two editions of the *Forbes Stock Market Course*, a comprehensive guide to investing and wealth management, and has authored articles for leading peer-reviewed journals, including the *Financial Analysts Journal, Journal of Financial Research* and *Financial Services Review*. Dr. Horan's award-winning research has been profiled in leading practitioner publications, including *The CFA Digest* and *Barron's*, as well as the popular press. He has received six research grants, including two from the Research Foundation of CFA Institute, and six research awards. He is Associate Editor of the *Financial Services Review* and has served on the editorial board of *The CFA Digest* as well as Director of Education for the CFA Society of Buffalo and several community boards.

Stephen earned a BBA in Finance with a minor in Mathematics from St. Bonaventure University and a Ph.D. in Finance with a minor in Economics from State University of New York at Buffalo.

Author:
Stephen M. Horan, Head, Private Wealth
CFA Institute
560 Ray C. Hunt Drive
P.O. Box 3668
Charlottesville, VA 22903-0668
US
Tel: +1 434 951 5694
Fax: +1 434 951 5370
Email: stephen.horan@cfainstitute.org