INVESTMENTS
Lecture 6: Some Cutting Edge Topics

Philip H. Dybvig
Washington University in Saint Louis

- Taxes
- Transaction costs
- Spending rules

Taxes

The optimal portfolio strategy to follow in the presence of taxes is one of the most challenging aspects of investments. It is also one of the most rewarding, since the potential gain from saving taxes is very large and often without risk. Earning 200 basis points on average above the market is not so easy; good decisions on taxes may earn 300 to 500 basis points without any accompanying risk.

Tax arbitrage

The purest form of tax strategy is a tax arbitrage. It is the job of the finance professional to identify tax arbitrage; it is the job of the legislature and the IRS to eliminate arbitrages that have been found! As an example, suppose two riskless investments have different amounts of taxable income associated with them. Then a high tax rate individual will go long the less taxed instrument and short the other, while a low tax rate individual will take the opposite position. This has the net effect of transferring taxable income from that high bracket individual to the low bracket individual, which results in a net savings of tax for the two. Given that the two riskless assets have different tax treatment, there is always a range of relative prices for the two assets which makes the trade profitable to both individuals.

In-class exercise: tax arbitrage

Consider two riskless investments over a year, “A” which is taxed at the ordinary income rate and bears interest of 10%, and “B” which is taxed at 1/2 the ordinary income rate and bears interest of 8%. Assume short sales are allowed with symmetric tax treatment. Is there a tax arbitrage for an investor in a 30% tax bracket? for an investor in a 40% tax bracket?

Impediments to tax arbitrage

Naturally, tax arbitrage is not popular with the IRS and legislators! One way they limit the availability of tax arbitrage is by giving short positions a different tax treatment than long positions. For example, a taxable individual cannot short a municipal bond ("muni"). And, capital losses can typically offset capital gains but not ordinary gains.

Tax dominance

Even when arbitrage is not available because of the tax treatment of short sales, we can talk about one investment dominating another from the perspective of a particular investor.

For a par coupon bond, the after-tax cash flows usually look like a par coupon bond. For an investor with tax rate t, the after-tax yield is the stated coupon yield times $1-t$. For discount or premium bonds, matters are more complicated and the exact tax rules for amortizing gains come into play. In the US, the code for amortizing gains is generally consistent with a constant interest rate equal to the yield of the bond.

Break-even tax rate

It is often useful to think about the critical "break-even" tax rate at which two investments are equally attractive. For example, consider an investor choosing to buy one of two bonds, a Treasury bond or an insured municipal bond (which we will treat as riskless), to be held to maturity. They are both being issued at par. The Treasury bond is taxable and bears interest at the rate of 5%, and the muni bears interest at the rate of 4%. For an investor with tax rate t, the after-tax rate-of-return on the Treasury is $5\%(1-t)$ and the after-tax rate-of-return on the muni is 4%. In this example 20% (the solution to $4\% = 5\%(1-t)$) is the break-even tax rate.

In practice, the break-even tax rate for munis is surprisingly low. Perhaps that is because few individuals want to invest in such an illiquid investment; part of the illiquidity is due to the fact that while interest on munis is not taxable, capital gains are.

In-class exercise: break-even tax rate

Suppose at a maturity of 30 years, the par Treasury rate is 8% and the par muni rate is 6%. What is the break-even tax rate?

Tax timing options

In the presence of taxes, holding an asset to maturity may be much less valuable than selectively choosing whether to the sell the asset depending on whether its value has risen or fallen. The value of the tax savings from careful timing of sales has been referred to as the value of a "tax timing option." Option pricing theory has been used to evaluate the tax timing option in some special examples, although so far we have needed to make strong and unreasonable assumptions to perform the analysis. In practice, valuation is complicated by the fact that the value of realizing losses depends on the tax status of the rest of the portfolio (e.g., whether there are capital gains that would not otherwise be offset).

General tax strategies

In general, optimal investment strategies in the presence of taxes are very complicated and difficult to compute. (For example, see my paper with Koo on my research page.) Here are two rules of thumb for taxable investors:

1. It is a good idea to realize capital losses whenever you can use them. If you are a type of investor that has unlimited loss carryforward, it may be a good idea to realize losses for future use.

2. It is a bad idea to realize capital gains unless it is necessary to realize gains to avoid an unacceptably high risk exposure.

Performance measurement for taxable accounts

Performance measurement for taxable accounts is a tricky matter, especially when different managers are being evaluated in different time periods. It is typical of optimal tax strategies that taxes are deferred, not avoided, and that makes it difficult to choose how to allocate gains to different periods. Several methods of accounting for performance of taxable accounts are used in practice, but to my mind none are completely satisfactory. I would like to work on the problem of devising improved performance measures for taxable accounts.

Wash sales

One issue in taxable investments is the so-called "wash sale" rule that says that you cannot deduct a loss if the position is reinstated within a month, in which case the two events are netted out for tax purposes. This rule is intended to minimize sales motivated primarily for tax timing. The rule is not very effective since there are many close substitutes in financial markets and therefore it is easy to reestablish an economically similar situation without invoking the wash sale rule.

Accumulating tax-free

The favorable tax treatment of some individual retirement plans is in the form of allowing the funds to accumulate tax-free with taxation of gains only at maturity. Another favorable treatment is if the income generating the initial investment is not taxed going in but is only taxed at maturity. Assuming a constant tax rate t (over time and for income as well as gains), we have that an investment of $\$100$ of pre-tax income at a riskless interest rate r for T years will grow to $\$100(1-t)(1+r-tr)^T$ without special tax treatment, to $\$100(1-t)+(\$100(1-t)(1+r)^T-\$100(1-t))(1-t)$ with tax-free accumulation, and to $\$100(1+r)^T(1-t)$ if taxes are paid only at the end.

Tax-free accumulation: numbers

The following table illustrates the advantage of tax-free accumulation assuming an interest rate of $r = 5\%$ and a tax rate of $t = 30\%$. The differences become more dramatic the higher the interest rate and the higher the tax rate.

# years	pay as you go	pay in and out	pay out only
1	72.45	72.45	73.50
5	83.14	83.54	89.34
10	98.74	100.82	114.02
30	196.476	232.77	302.54

Transaction costs

Transaction costs represent an important but incompletely understood component of investment performance. Here are some thoughts:

- Transaction costs imply that we will stray somewhat from the ideal mix of assets (either across or within asset classes) before we will trade.
- As costs increase from zero, it is optimal to make the no-trade range large to avoid transaction costs which are of first order of importance rather than maintain risk exposure which is of second order.
- It is useful to think of costs as being spread over the holding period of an asset. A bid-ask spread of $\$0.25$ on a stock costing $\$50$ and held for 2 years can be thought of as reducing the return by about $\$0.25/(\$50 * 2) = .0025$ or 25 basis points. This approximation makes the most sense over horizons much shorter than the duration of the investment itself.
- Most investment managers trade more than can be justified by benefits of trading to the clients, in part because they have to meet benchmarks.

In-class exercise: transaction costs

Consider two stocks each priced around $\$50/share$. One has a spread of $\$1/8$ and the other has a spread of $\$3$. What is the spread's adjustment to expected return when the investment horizon is 1 year? when the investment horizon is 20 years?

Round lots

In each exchange, there is a minimum number of shares, called a round lot, that can be purchased without facing larger commissions and spreads. In general, we prefer to trade in round lots to avoid the higher costs. For smaller amounts of money under investment, this effectively limits the amount of diversification we can do (unless we buy some sort of mutual fund or index product).

Execution

There are lots of ways to execute trades. For example, for NYSE stocks, we can send a market order or a limit order to the NYSE. Or, we can "work" the order, trying to execute our trade in pieces over time when we believe we can get the best price. Alternatively, we can send our order to an electronic clearing network that will execute the trade overnight or during the day in an auction with other big institutional investors. Unfortunately, there is little hard evidence on how we can minimize execution costs, and it is hard to infer this information from available price and quantity data. It makes economic sense that there should be some trade-off between immediacy (or what amounts to the same thing, probability of execution) and price, but we do not know even this for sure. It does seem obvious that the spread and quality of execution is potentially the largest part of transaction costs, especially for thinly traded stocks.

Spending rules and asset allocation

The remainder of this lecture is devoted to spending rules and their connection to asset allocation. This is connected closely to my work that was awarded the Common Fund Prize. That prize was awarded for my paper "Duesenberry's Ratcheting of Consumption..." in recognition of its importance for the practice of managing university endowments.

Repeated single-period perspective

We can usefully think of the portfolio management problem as an alternating sequence of choices of asset allocation and spending rule. At the start of an investment period, we have a certain amount of endowment to allocate to investments in different assets. At the end of the period, our assets have grown or declined in value, due to our skill and luck. At that point, we decide how much of our current endowment to include in this year's budget and how much to reinvest. The outcome of that decision gives us the start-of-period endowment for the next asset allocation problem.

The single-period perspective treats each asset allocation or spending decision as a separate choice, linked only through the fact that previous decisions determine what resources are now available.

Single-period perspective: limitations

1. asset considerations
   - Commitments to illiquid assets (e.g. real estate or private placements) must be made over many periods.
   - Transaction costs are incurred in a single period but may be justified by benefits over time.
2. preference considerations
   - It is difficult to assess the value of returns over a single period without knowing the pattern of use of funds in the future.
   - (my main point) We have different preferences about the maintenance of existing programs and commitments than we do about starting new projects.

Examples of Common Practices

1. Spending rule
   - Plan to preserve capital.
   - Spend a percentage based on some historical return on capital.
   - Using a moving average to smooth income (partial recognition of my main point, but not the best response).
   - Whether the percentage is done on a unit basis is important when new money comes in.
2. Investment policy
   - Ranges for proportions are chosen.
   - Investment choices depend little on the history.
   - If anything, losses are followed by a more aggressive policy to try to make up lost ground.

Features of my proposal (simple extreme case)

1. Spending rule
   - Plan to preserve capital to an extent dictated by time and risk preferences.
   - Increase spending on essential activities only when the endowment reaches a new maximum.
2. Investment policy
   - Conceptually, separate out the portfolio into a committed account and a discretionary account.
   - The committed account is immunized.
   - The discretionary account is invested in fixed proportions.
   - When spending increases, the funds needed to ensure committed spending are transferred from the discretionary account to the committed account.
   - The investment policy is consistent with the CPPI analyzed by Black and Perold.

More on my Proposal

1. General features within the model
   - declining or rising commitments over time (e.g. attrition in faculty)
   - some expenditures not committed at all in advance: funded by a separate account with proportional spending
2. General features requiring adjustment of the model
   - stochastic interest rates
   - commitment to an irregular pattern of commitments
   - reliance on other sources of income (e.g. tuition)

What is committed expenditure?

Useful models are simplifications of the world. In practice, commitments are not absolute. For example, we can actually eliminate valued departments or buy off tenured faculty for less than the present value of their future salary. Or, we can possibly sell our new building if we are unable to meet the mortgage payments. However, these are desperate actions and it is a good approximation to think that we want to be certain to make these expenditures.

Basic assumptions

1. Investment opportunities are the same at different points in time. A riskless asset bears a constant interest rate, and the mean and standard deviation per unit time of the efficient risky portfolio is constant.
2. Consumption (or more generally consumption of some goods) is not permitted to decline (more general at more than some rate).
3. Preferences are the same looking forward from any point in time, given consumption does not decrease. Risk and time preferences are scale independent.
4. The model is similar to the Merton model except for the constraint on non-declining consumption.

Performance comparison with traditional strategies

1. qualitative properties similar to portfolio insurance in some ways, but without the jerky restart every year
2. much better worst-case scenario (persistent declining market)
3. better best-case scenario
4. worse in up-and-down markets (whipsaw effect)