

The “Variety of Security Returns” and the Tradeoffs Between Capital Growth, Risk and Taxes

Dan diBartolomeo
High Net Worth Seminar
January 2008

The Problem and an Assertion

- Our experience suggests that asset managers are unable to parameterize portfolio construction problems so as to provide an optimal balance between capital growth, risk control and tax avoidance for taxable investors
- We assert that the key to correctly balancing tax avoidance with risk adjusted returns is understanding the economic value of the option to realize capital gains at a time of the investor's choosing
- The value of the "tax timing" option is directly related to the expected cross-sectional dispersion ("variety") of asset returns within the investor's portfolio

This Presentation in Outline

- Describe two modifications of the traditional mean-variance utility function appropriate to taxable investors, augmenting the objective function
- Incorporate a “structural” form of estimation error that has been relatively less explored in the finance literature
 - the “single period” assumption built into the Markowitz mean-variance optimization process
- Show how the cross-sectional dispersion (“variety”) of security returns within a market is the crucial determinant of “effective” tax rates
- Using large scale numerical simulations, explore the empirical relationship between “portfolio variety” and maximum tax efficiency
- Incorporate our simulation results into asset allocation and portfolio optimization taking tax efficiency into proper account

We Walk From Where We Stand

- Markowitz and Levy (1979) propose a function of mean variance as a representation of investor utility

$$U = R - S^2 / RAP$$

- This is just risk adjusted return where the size of the risk penalty can be scaled to fit the investor's risk tolerance
- Assumptions
 - The parameters of return distributions are known with certainty. In the real world, we can get it wrong.
 - The future is one long period in which our input parameters values (that are both correct and certain) never change.

Case 1: Taxable Asset Allocation

- As described in Wilcox (2003), the key issue in formulating investment policies is how aggressive or conservative an investor should be to maximize their long term wealth subject to a shortfall constraint (a floor on wealth). One way to express this for a taxable investors is:

$$U = E\{ R * (1-T^*) - L S^2 (1-T^*)^2 / 2 \}$$

- L is the ratio of total assets/net worth from a life balance sheet
 - In Northfield terminology $RAP = 2/L$
 - T^* is the effective tax rate
- Today's first empirical question is how to estimate the possible differences in T^* , the effective tax rates across different asset classes

Case 2: Taxable Portfolio Management

- Under US tax law, we can realize capital gains on a “lot by lot” basis. In such cases we can model capital gain taxes as explicit transaction costs, rather than as a scalar on returns

$$U = E\{ R * - L S^2 (1-T^*)^2 / 2 \} - (C * A)$$

C = expected “transaction costs” including CG taxes

A = rate at which costs are amortized over the economic life of an event (reciprocal of the expected holding period)

- We need differentiate between trading costs and capital gain tax, but how?
 - If we don’t trade, we don’t have a trading cost
 - If we don’t realize an embedded capital gain now, we’ll probably realize it sometime in the future, owing tax then
 - Cost avoidance, but tax deferral

Initial Thoughts on Single Period Assumption

- In the real world, things change and our parameter estimates for return and risk (even if initially exactly correct) are likely to change as well.
- If transaction costs are zero, we can simply adjust our portfolio composition to optimally reflect our new beliefs whenever they change.
- If transaction costs are not free, the single period assumption is a serious problem.
- If transaction costs are large (e.g. capital gain taxes), the single period assumption is wholly unrealistic. Tax authorities also seem to be interested in things like weeks, months and especially "tax years."

Geometric Versus Linear Tradeoffs

- For small transaction costs, arithmetic amortization is sufficient, but if costs are large we need to consider compounding
- Assume a trade with 20% trading cost and an expected holding period of one year.
 - We can get an expected alpha improvement of 20%. But if we give up 20% of our money now, and invest at 20%, we only end up with 96% of the money we have now.
- Solution is to adjust the amortization rate (A) to reflect the correct geometric rate

$$A_g = 100/(100-A)$$

- Our estimation of T^* as the “effective tax rate” will implicitly include this issue

A Simple Rule for Better Tradeoffs

- Even if we are amortizing our costs sensibly, we are still maximizing the objective function to directly trade a unit of risk adjusted return for a unit of amortized cost per unit time.
 - This is only appropriate if we are certain to realize the economic benefit of the improvement in risk adjusted return, which is only true over an infinite time horizon
 - We propose to adjust the amortization rate to reflect the probability of actually realizing the improvement in utility over the expected time horizon, and the investor's aversion to the uncertainty of realization

$$U = R - S^2 / RAP - (C \times \Gamma)$$

$$\Gamma = A_g / (1 - M * (1-P)), M = 1 - (RAP/200)$$

P is the probability of realizing the improvement in risk adjusted return over the expected time horizon and L is the range of (0,1)

We Live in a Multi-period World

- If our portfolios have finite holding periods rather than the infinite single period assumed in Markowitz, we have a probability between .5 and 1 of actually realizing an expected gain in utility
- We define the probability of realization, P , like a one-tailed T test

$$P = N\left(\frac{(U_o - U_i)}{TE_{i0}} * (1/A)^{.5}\right)$$

$N(x)$ is the cumulative normal function:

$$N(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-u^2/2} du$$

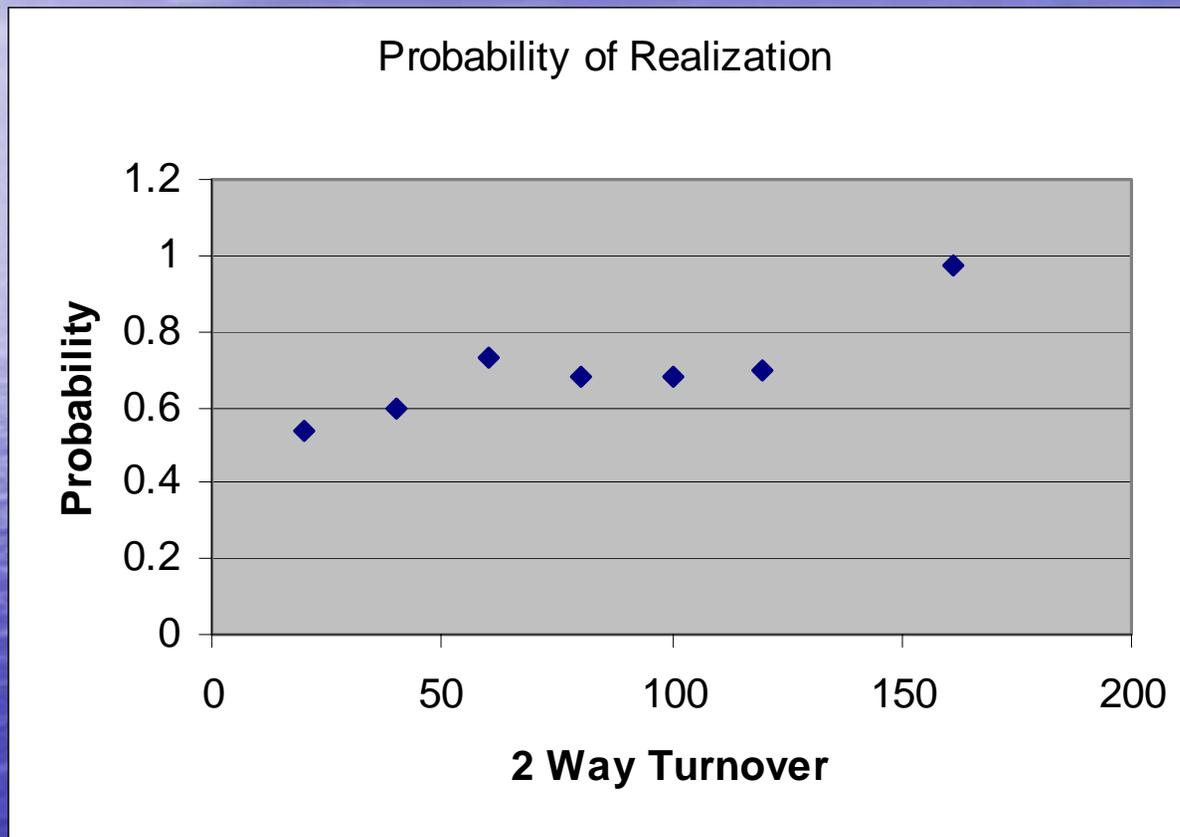
The Realization Probability

- The numerator is the improvement in risk adjusted return between the optimal and initial portfolios
- The denominator is the tracking error between the optimal and initial portfolios. Essentially it's the standard error on the expected improvement in utility
 - If there is no tracking error between the initial and optimal portfolios, P approaches 100%. Consider "optimizing a portfolio" by getting the manager to cut fees. The improvement in utility is certain no matter how short the time horizon.
 - *Not something to which we usually pay attention*
- If turnover is very low, A will approach zero, so P will approach 100%. For long time horizons, we have the classical case that assumes certainty

Empirical Example

- Initial Portfolio 85 Large Stocks
- S&P 500 Benchmark
- Random Z-scores as alphas
- Risk Tolerance = 40
- Expected Turnover 25% per annum
- No position bigger than 3%
- No position smaller than .25%
- 20 cents per share trading costs

Probability of Realization



Valuing the Tax Timing Option

- US tax law provides for selective realization of capital gains at the tax lot level. Investors have the option of which gains or losses to realize in their portfolio.
 - If all the positions in your portfolio have the same degree of capital gain, you are indifferent as to which gain you realize. For the tax timing option to have value, you must have dispersion in the degree of percentage capital gains across the tax lots of the securities in your portfolio.
- You can get dispersion in degree of capital gain by either
 - Reinvesting income over time at different prices into the same asset (Horvitz and Wilcox, 2003)
 - Owning multiple securities that have dispersed returns. The larger the dispersion, the larger the opportunity set for intelligent offsetting of gains and losses to defer tax

Estimating the Relationship of Variety to Tax Efficiency via Monte Carlo

- Start with a data set of stock returns
 - Monthly return time series for the 850 stocks in our database that had no CUSIP changes (no major corporate actions) from 1990 through 2004
 - Convert the returns to cross-sectional Z-scores
- You can now construct simulated returns for every security, for any chosen values for the expected mean and cross-sectional standard deviation
 - Preserves the correlation structure across stocks
 - Preserves the relationship between market volatility and cross-sectional volatility
- Using “bootstrap” re-sampling methods, we can construct as many simulated return histories of any length

Simulate Portfolio Tax Behavior

- Start with an equal weighted portfolio of N securities
- Assume some expected return on the market, with a fixed dividend yield
 - Reinvest dividends in an ETF with zero dividend yield
- Assume some amount of monthly turnover associated with active management to account for transaction costs
- Pick a time horizon (e.g. 25 years) at which time the investor is assumed to die
- Roll the positions forward month by month until the time horizon

Measuring Effective Tax Rates

- For each simulated series of events, we calculate three rates of return
 - The total return on the portfolio assuming zero capital gain tax
 - The after tax total return assuming capital gain tax via an annual “mark to market”
 - The after tax total return assuming we use our assigned degree of turnover to sell positions that have the biggest losses or smallest gains to defer taxes as long as possible
 - Due to tax deferral, effective tax rates will be much lower than nominal capital gain tax rates
- The differences in these after-tax returns represent our ability to manage realization of capital gain taxes
- Selling losers and keeping winners will tend to concentrate the portfolio over time, increasing risk
 - We reinvest dividends into an ETF as a simple way to create an offsetting degree of diversification

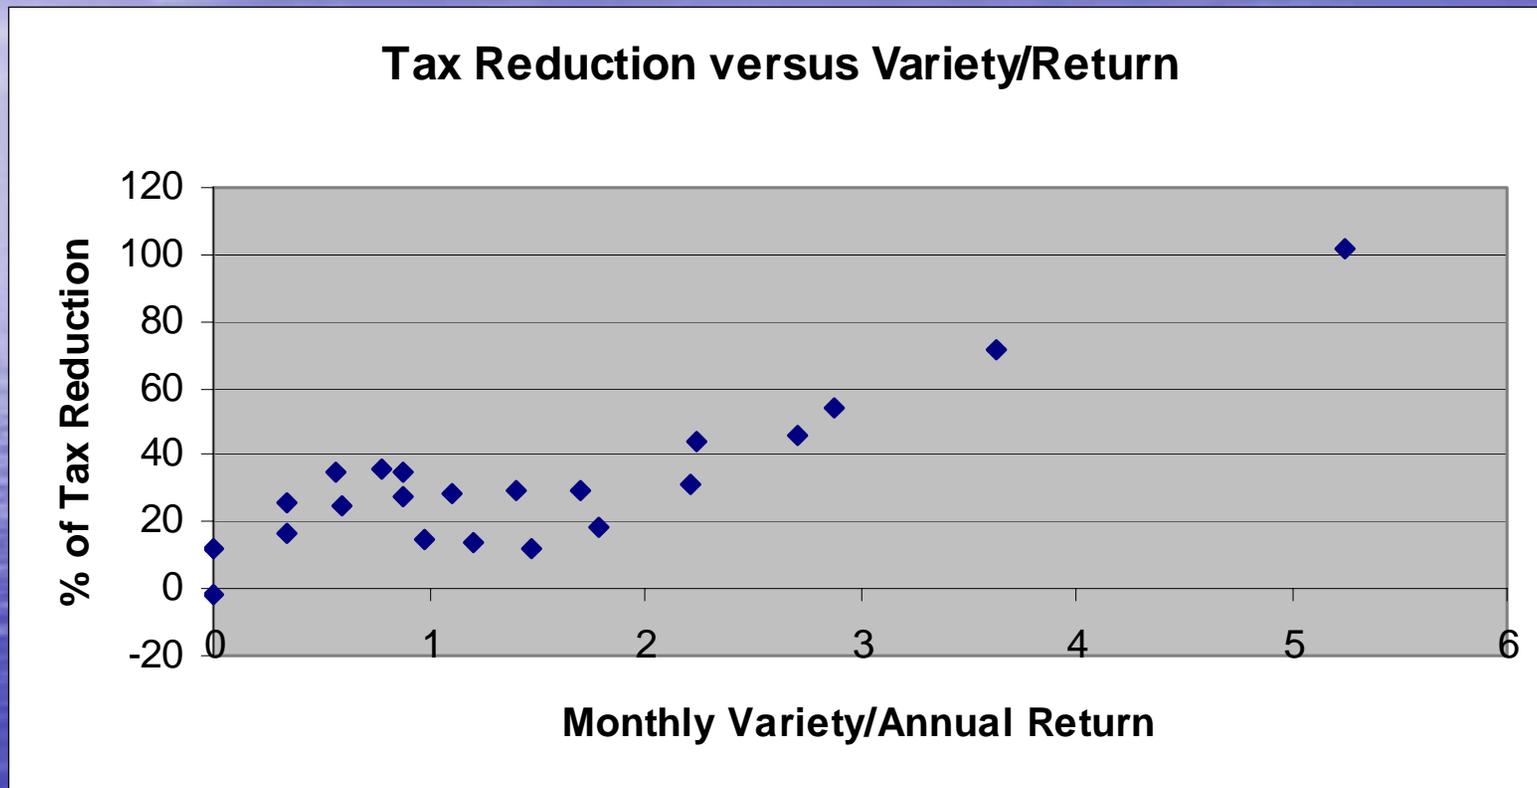
Valuing the Tax Timing Option

- We ran several hundred portfolio simulations at each possible combination of four parameters
 - Initial number of securities from 20 to 100
 - Annual turnover from 0 to 100%
 - Monthly cross-sectional dispersion from 0 to 30%
 - Time horizons from 24 to 600 months
- Typical parameters
 - Expected market return 9%, with dividend yield of 2%
 - Nominal capital gain tax of 15%, income tax 40%
 - .5% round trip trading costs
- Calculate the net effective tax rate for each simulation and average across the sample at each level of cross-sectional dispersion

Estimate a Relationship

- Measure the opportunity to reduce taxes through selective realization as the ratio of cross-sectional dispersion to the expected market return
- Define “reduction in taxes” as the fraction of taxes saved by selection tax realization as compared to “market to market” where the investor has no timing option
- Effective tax rates of 12% down to .8% (zero CG tax)
- Over sample of approximately 500,000 portfolio simulations grouped into twenty two batches
 - Correlation in the hypothesized relationship is .88, which is barely statistically different from one

Variety and Tax Efficiency



Empirical Highlights

- For cross-sectional dispersion values typical of US large cap stocks (i.e. 10% monthly), we get effective tax rates around 5 or 6%, a lot lower than the nominal tax rate
- For US small cap or international stocks, dispersion values of 25 to 30 are typical, leading to effective tax rates below 3%
 - At average turnover levels of 50% per annum with a 25 year time horizon, we often have enough tax losses to stay in a net negative realized capital gain situation indefinitely, if we use all our turnover to maximize the tax timing option
- We can now get an estimate of T^* for an chosen combination of our parameters such as cross-sectional dispersion, tax rate, market return, turnover, etc.

Putting Simulation Output to Work

- For asset allocation problems, we can estimate T^* for each asset class we are considering and plug directly into our Markowitz asset allocation framework
- For tax sensitive active management:
 - Estimate the effective tax rate under the maximum use of the tax timing option. This is the minimum tax we can expect to pay given our assumed turnover
 - To the extent that this minimum is above the effective tax rate for passive management, it is a reduction in portfolio alpha. Essentially an ongoing cost like a management fee
 - The excess of the nominal capital gain tax rate above the minimum effective capital gain tax rate becomes the capital gain tax rate that is included with transaction costs. It's the "optional" portion of tax that we choose to tradeoff against expected improvements in alpha or reductions in risk

Conclusions

- Portfolio construction for taxable investors requires attention to taxation at both the asset allocation and portfolio trading levels
- The single period assumption in MVO implies that trading costs and improvements in utility can be traded as if both are certain
 - Finite holding periods imply that the improvement in utility is uncertain. We must therefore consider the probability of realizing an improvement in utility arising from an optimization as being between 50% and 100%
- Simulation methods can be used to estimate effective tax rates for wide range of market conditions and portfolio management practices
- We have shown that the opportunity to manage taxes effectively is a roughly linear function of cross-sectional dispersion of security returns

References

- Levy, H. and H. M. Markowitz. "Approximating Expected Utility By A Function Of Mean And Variance," *American Economic Review*, 1979, v69(3), 308-317.
- Wilcox, Jarrod. "Harry Markowitz and the Discretionary Wealth Hypothesis", *Journal of Portfolio Management*, 2003.
- Horvitz, Jeffrey and Jarrod Wilcox. "Know When to Hold 'em and Know When to Fold 'em", *Journal of Wealth Management*, 2003.
- diBartolomeo, Dan. "Applications of Portfolio Variety", in *Forecasting Volatility*, ed. Steven Satchell and John Knight. Butterworth-Heinemann, 2005. Oxford.