NORTHFIELD modeling financial markets worldwide

Generating Tax "Alpha" for Private Wealth Households

Dan diBartolomeo

Northfield Webinar

April 10, 2014

Goals for Today's Discussion

- Describe analytical processes for taxable portfolio management that can add significant amount of tax alpha to passive or active strategies
- Providing an intuitive decomposition of the portfolio optimization process
- Examine the huge tax alpha associated with resolution of concentrated positions
- Review the process of being tax efficient in a multiple-manager account structure
- An empirical study of tax timing options and expected value of resultant tax alpha

Why is Tax Alpha Important?

- With the traditional concept of alpha in asset management, every active manager would like to believe that they can outperform, but it's mathematically impossible. For some to be above average, another manager has to be below average. Alpha is a scarce good and very difficult to consistently obtain.
- With tax alpha, *all managers can improve private client outcomes in a non-competitive and reliable process.* Clients win with better after-tax returns and all managers can win by improving their client outcomes. The only loser is the tax collection authorities, who can then foster legislation to change the rules of the game.

"Manufacturing" Tax Alpha

- "Mass customization" is the key to managing large volumes of high net worth clients. Separate the investment process into three distinct roles
  - Client relationship to define client needs and wants
  - Investment research and "best ideas" investment models
  - Adaptation of "best idea models" to individual client needs. This is the usual bottleneck for large volumes, but it can now be effectively automated
- Most firms do a poor job with private client accounts
  - Taxes are often ignored, or treated as a last minute afterthought
  - Or they use taxes as an excuse to do "pseudo passive" management at active management fees

Where Most "Tax Sensitive" Firms Fall Short

- The most common failure is to treat security portfolios across asset classes in a unified fashion for taxes
  - A capital loss on a bond can be used to offset a gain on a stock
- Rule based "tax sensitive" processes are flawed
  - Many systems allow you to program in rules like "don't sell a tax lot at a short term gain if it has been held 10 months or more"
  - Such a rule would has been a disaster in a case like Enron
- Managing taxes as a constraint is almost as bad
  - Some firms operate as "maximize return subject to not taking more than $X in net gains during this tax year"
  - The value of X is usually arbitrary and does not reflect rational tradeoffs between return, risk, current taxes and deferred taxation

The Tax Alpha Assembly Line

- Use "tax aware optimization" to index track the "best ideas" model portfolios to different clients
  - Adjust alphas to reflect expected after-tax dividend stream
  - Treat capital gain taxes as big transaction costs
  - Amortize taxes to reflect portfolio turnover, compounding value of tax deferral, and likelihood of stepped up cost basis
  - Adjust risk aversion to both absolute and index relative risks
  - Adapt number of stock names to different portfolio sizes
  - Exploring the tax/risk efficient frontier to maximize long term wealth accumulation for transition of legacy portfolios
- Emphasize tax aware strategies
  - Value strategies sell what went up. Bad for taxes compared to momentum strategies at the same turnover
  - Quant strategies are more flexible than bottom up stock picks
  - Stock universes with higher cross-sectional dispersion (e.g. small cap) increase the value of the tax deferral option

Optimization Objective Function

- Investor objective is to maximize risk- adjusted returns net of costs:

$$U = \alpha - ((\sigma_s^2 / RAP_s) - (\sigma_u^2 / RAP_u)) - ((C + T) * A)$$

Return | Risk Component | Implementation Cost

  - Most commercial portfolio optimizers uses the objective function described in Levy and Markowitz (1979).
  - Portfolio return variance is the proper measure of risk because the difference between the arithmetic average rate of return and the geometric average rate of return is proportional to the variance (see Messmore, 1995)

Objective Function Explained

- $U = \alpha - (\sigma_s^2 / RAP_s) - (\sigma_u^2 / RAP_u) - ((C + T) * A)$

  - **α = the "certainty equivalent" expected portfolio return**
  - σ_s^2 = portfolio variance risk due to common factors (correlation across securities)
  - σ_u^2 = portfolio variance risk due to stock specific risks
  - **RAP = risk tolerance**
  - C = transaction costs for the optimization
  - T = capital gain taxes for the optimization
  - **A = amortization constant**

Items in **bold** are what user's have responsibility to set sensibly.

$RAP = 1/\lambda$

Cost Amortization Constants

- A good starting value for the amortization constant is the percentage expected annual "one way" turnover
- Adjust for:
  - Time value of money. Deferring costs is more valuable when rates of return are high
  - Possibility of stepped up basis at death. If you're 99 years old, you don't want to pay any capital gains now
  - The incremental tax on short term gains over long term gains is always spread over the "remaining time to long term status"
- In 2009, Northfield introduced a change in the amortization constant process that approximately corrects for the traditional Markowitz assumption that the future is one long period.
  - Performance measurement services and the IRS don't agree on how to treat taxes
  - Be careful of fiscal years that that are not calendar year
  - http://www.northinfo.com/Documents/500.pdf

An Intuitive Explanation of Tax Optimization

- Assume we can create portfolios that we feel would be optimal.
- If taxes and other constraints were not in the way, we could simply trade to switch to the optimal portfolio from the current portfolio
  - We can describe the distance between where we are and where we want to be in terms of tracking error.
- The list of required trades will divide into three categories
  - Trades that reduce tracking error and generate tax losses
  - Trades that reduce tracking error and generate tax gains sufficient to offset the losses we just got
  - Trades that further reduce tracking error but cost us something in net realized capital gains
  - We should always do the first two sets of trades
  - What portion of the trades in the third set we should do depends on our willingness to trade taxes for expected return and risk

Dealing with Concentrated Positions

- Real life case:
  - What do you do with a client that has $600 Million dollars in one stock with near zero cost basis?
  - The stock is traded on the NYSE
  - According to the risk models, the stock has an expected volatility of about 50%
- Truncate the risk using a "costless collar"
  - Write call, use the call proceeds buy put
  - Option transaction costs can be high
  - It's a holding action at best, doesn't solve the problem
- Prepaid variable forward
  - A fancier version of the collar strategy using an OTC derivative
  - Typically very expensive to do a custom OTC deal with an investment bank, and you have built-in counterparty risk

Working out of a Concentrated Position

- A leveraged complementary fund is the way out
  - Margin the big holding to get fresh cash
  - Short an index ETF to bring market exposure down
  - Use the fresh cash to fund a complementary fund that offsets the factor risks of the big position using a mix of volatile stocks
  - Harvest taxes losses in the complementary fund, and use to offset gains on sales from the big position. The big position gets whittled down over time
  - The risk is a having to liquidate during a broad market decline that creates a short term capital gain on the ETF
- Over two years it was possible to reduce the big position from 100% of the portfolio to about 25% with zero net realized capital gain
  - Portfolio volatility was cut in half for a minimum tax alpha of over 1000 basis point per annum through improved compounding of the less volatile returns

CPM and The Best Wine We Ever Served

- Centenary Solera Madeira, 1845 was served at our 1992 client conference at Bretton Woods
  - This was the last time I spoke about Centralized Portfolio Management at a Northfield event
  - Brokers thought it would reduce trading (and their incomes)
  - Asset managers thought it would cause reductions in management fees (and hence their incomes)
  - Plan sponsors thought I was calling them stupid for not having done it already
  - Everyone agreed it was probably analytically correct but were very unhappy with me. As usual, I sought help from the bartender. The Madera was hastily added to the dinner menu.
- Everyone, including me, agreed that no one would ever do it in the real world. *Everyone was wrong.*

Private Clients Using Multiple Managers

- *The purpose of using multiple active managers is to add return*, not to provide diversification.
  - If an investor wishes to simply reduce the risk of under-performing a particular benchmark index, it may be accomplished efficiently by owning an index fund. Such passive investment also reduces fees and trading costs
  - It must therefore be true that the value added by active management arises from the manager's efforts to create forecasts that are meaningfully predictive of future index-relative security returns (that is, alpha)
  - There may be some logic in having multiple active managers in an effort to prudently protect the fund against embezzlement or other illegal actions on the part of investment agents, but this protection could be met equally well with multiple passive managers

The Theoretical Multi-Manager Problem

- In the typical multiple manager portfolio, a single manager has influence over each dollar of the total fund.
  - This methodology rests on two beliefs: that each manager is expected to produce superior performance on average and that the certainty of any particular manager producing superior performance is low.
  - Hence, there is a need to diversify with multiple managers so as to decrease the uncertainty associated with each manager's forecasts of future alphas.
  - By allowing multiple managers to operate within the same universe of securities with similar constraints, the aggregate position of the total fund with respect to any particular security is reflective of the consensus alpha forecast of the managers. In this fashion, the fund benefits from the improvement in the aggregate quality of the return forecasts.

Multi-Manager Correlation

- Unfortunately, there exists the possibility that manager alphas will be correlated.
  - For example, if a fund had ten managers, but all ten simultaneously adopted a similar strategy (e.g., mid-cap, low P/E), little improvement in the quality of the alpha forecasts would result.
  - Even worse, the overall fund would be subject to large residual risk. The presumed diversification benefit of multiple managers would vanish if each manager held a similar portfolio
  - To prevent such situations, it is common practice for plan investors to hire multiple active managers, but with distinct mandates. For example, the investor would hire separate managers for large-capitalization and small capitalization stocks (or growth/value, etc.).
  - Once we have separated the manager mandates into mutually exclusive habitats, we now have only the predictive skills of a single manager applying to each security. *The improvement in aggregate forecast quality has been lost.*

The Crossing Cost Problem

- One way to preserve the benefits of multiple managers without using distinct habitats would be to allow all managers the same mandate, but select managers that will be uncorrelated in their future alpha forecasts.
  - This unfortunately leads to the problem of internal crossing.
- Let's assume that the price of stock XYZ has just run up in reaction to the announcement of better than expected earnings. Value manager "A" believes that stock is now overpriced and sells it. Growth manager "B" believes that the stock has good future prospects and buys it.
  - The investor has just incurred two transaction costs, *possibly a capital gain tax, a possible wash sales violation,* and continues to pay two active management fees while nothing has happened to the total portfolio.
  - When we expand this example to industries and sectors, multiple active managers will inevitably spend most of their time in a meaningless but expensive tug-of-war.

Capital Allocation: Multi-manager Portfolios

- Allocating capital across multiple active/passive managers is a complex issue
  - Managers can be of varying levels of aggressiveness
  - Active managers generally are judged on *pre-tax* benchmark relative performance, while investors receive after tax returns
  - Investors have varying degrees of concern about both benchmark relative and absolute risk
  - The correlation of active returns across managers can be very unstable as market conditions change
- Shifting money between managers to rebalance capital allocations is expensive
  - *May involve capital gain taxes as well as transaction costs*
  - *Managers are hesitant to accept "legacy" positions from other managers that could reduce their performance record*

Multi-Disciplinary Accounts Distort Intent

Assume we have a two manager portfolio with 50% of capital allocated to each.

- Manager A thinks XOM is a great stock and overweights it 3%. He has market weighting on the airline sector. Manager B has a big 5% overweight in airlines (which he considers undervalued) and so overweights XOM 3% as a hedge against oil prices rising.
- If we just add the positions together, we will have a 3% overweight in XOM and a 2.5% overweight in airlines.
- However, Manager B wanted to overweight XOM as a hedge against his big airline bet. Since the airline sector is now just overweighted by 2.5% in the aggregate portfolio, the correct overweight for XOM in aggregate is 2.25% not 3%
- The process is inefficient because each manager makes risk/reward decisions myopically without knowledge or formal concern as to what the overall portfolio looks like.

Key Insight

- Almost all asset managers use risk models from an outside vendor for portfolio analysis and construction
  - Suggests that managers think the risk models work pretty well
  - Managers see their "value added" in superior return forecasting
- If everyone roughly agrees on the covariance among securities, then we can infer manager "alpha" forecasts from the portfolio they choose to hold.
  - Active managers must think their portfolio is optimal or they would hold a different portfolio
  - Expected returns must offset marginal risks at optimal weights
  - Sharpe (1974), Fisher (1975), Black-Litterman (1991)
- If we know the aggressiveness level of a manager's strategy, we can obtain the alphas directly.
  - If not we can estimate risk tolerance from observed risk values
  - Alternatively, we can use the implied rank values and then map into the expected cross-section of returns

Easy Things to Do with Implied Alphas

- Build customized "separate accounts" to meet the heterogeneous needs of taxable clients from a model portfolio such as a successful mutual fund
  - Easy to create portfolio versions at different aggressiveness levels, portfolio sizes, SRI constraints or tax considerations
  - Model portfolios can be run with any strategy, fundamental or quantitative
  - Relative performance across versions makes clear to investors the tradeoffs they are experiencing between "best ideas" investment performance and the influence of customizing to their preferences and constraints
  - Rationally migrate portfolios from legacy positions toward the model portfolio *over time to minimize capital gain realization*

Private Client CPM

- Combine "implied alpha" methods with Rosenberg (1977) "centralized adviser" structure
- Each external manager runs a paper portfolio, as is done with managers that provide model portfolios for "wrap" programs, reporting trades to the investor's agent
- Use statistical methods to combine the implied alphas across all managers, and apply to running one portfolio
  - The consensus alpha is used on the entire value of the portfolio, so you *get the return benefit of "two heads are better than one". We get the sum of the manager's knowledge, not the average.*
  - See Johnson (1972), Rudd and Clasing (1982)
  - Risk control is internally consistent as the alphas are implied using the same model used for portfolio construction

Benefits of CPM for Private Clients

- Almost all the "crossing costs" disappear
- Taxes are a lot easier to manager. No more wash sales
- Rebalancing costs disappear as shifting managers adjust changing weights in the formation of consensus alphas
- Manager aggressiveness and turnover levels no longer impact the aggressiveness of the whole
  - They are just part of the weighting of the alpha consensus
- Managers that are capacity constrained can participate
- No more myopic risk controls at the individual manager level filter through to the central portfolio
- *Implemented by Vanguard Australia in 2007 with a 50-60 basis point alpha and 50% reductions in turnover*

Possible Pitfalls of CPM

- Implied alphas may be biased if inferred from long-only managers.
  - See diBartolomeo (2008) for a solution
- External managers could concentrate on illiquid positions to boost returns on their paper portfolio
  - Use an agreed upon market impact model to estimate "trading costs" of the external paper portfolios and judge performance after costs
  - We've provided an excellent market impact model free to clients since 2009
- Managers could "front run" the central portfolio if they know their weight in the consensus is high

T*, Tax Alpha and 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:
  - 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

The Relationship of "Variety" and Tax Efficiency

- Use Monte Carlo Simulation
  - Monthly return time series for the 850 stocks in our database that had 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 dispersion.
  - Using "bootstrap" re-sampling methods, we can construct as many simulated return histories of any desired 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 time 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 the T* to Variety 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 selective 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

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 of 25 to 30% are typical, with 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.
- *Reliable tax alpha of up to 100 basis points annually is demonstrable*

Conclusions

- All asset managers dealing with taxable clients should be making significant effort to capture the "non-competitive good" known as tax alpha
- In the case of portfolios with concentrated positions, the degree of reliable benefit from tax alpha is at least comparable in magnitude to the traditional alpha provided with far less certainty by even the best active managers
- Efficient processes have existed for more than twenty years to bring tax alpha to multiple manager structures
- *Simulation tests suggest that reliable tax alpha of up to 100 basis points annually should be routine.*