NORTHFIELD modeling financial markets worldwide

A Plausible Short Cut for Tax Optimization and Why Not to Use it

Dan diBartolomeo

Webinar
May 2023

Introduction

- Since the 1995 introduction of the Northfield taxable optimization process, the explicit objective was to be optimal in terms of risk-adjusted (see Markowitz and Levy, 1979) returns, after tax effects.
- In order to obtain such optimality, the effects of taxation must be included in the optimization objection function along with returns, risk, and trading costs.
  - This requires that rational economic tradeoffs be made between all components of investor utility in "certainty equivalent" form.
  - Alpha values (if any), risk, trading costs and taxes must be in the correct units and scaled appropriately.
- Some asset managers that market themselves as are tax efficient have chosen to operate with a much simpler objective, which we would argue is plausible, but inferior except in special cases.

The Big Short Cut in Words

- Rather than parameterize the full optimization objective function, we can express this idea as:
  - *Optimize the portfolio on a pre-tax basis, subject to a maximum dollar constraint on taxes to be paid during the current tax year.*
  - Alternatively, the constraint could be expressed as a maximum constraint on net capital gains to be realized during the current tax year.

- A purported benefit of this objective is that is far easier for a retail investor to understand in terms of "*we will do the best we can with your portfolio and try to limit what you pay in taxes this year on your portfolio profits to $X"* (presumably a small percentage of the portfolio value).

Why Does This *Seem* Like a Good Idea?

- It's simple, so retail investors understand what they can expect.
- The manager doesn't have to make any effort to select economically rational tradeoffs between taxes and risk
  - The reason this is complex issue is that taxes are paid by investors based on their absolute gains/losses, while the manager often chooses to define risk in benchmark relative terms (i.e. tracking error).
  - This is particularly true for passive asset managers who are not trying to add return, but merely minimize tracking error.
  - *Naive investors assume the effective tax rate on investing is their marginal tax rate, making the implicit "tax alpha" on a strategy with limited tax payments seem much higher than it really is.*
- Since index returns are not transparent on an after-tax basis, investors often judge passive managers on pre-tax returns.
  - A convenient approximate method for calculating after-tax benchmark returns was presented in Gulko (1999).

Why Is This Really Not a Good Idea?

- It's not clear that there is any economically rational method for a retail investor to set the "tax budget", as there is no attempt to tradeoff between return, risk and taxes.
  - Investors are averse to taxation and so may select too low a "tax budget", which has the *unfortunate effect of exaggerating the perceived tax alpha.*
- If the tax budget is not sensibly set it will be very hard to figure that out in terms of actual optimality.
  - The cost of typical optimization constraints (e.g. minimum dividend yield) can be estimated using the calculus method of "Lagrangian Multipliers."
    - See Scherer and Xu (JPM, 2007).
  - Capital gain tax effects are highly non-linear (short term, long term at different rates, wash sales, etc.) *so the impacts are neither continuous or differentiable.*
  - Asset managers are fiduciaries. They can't legally hide behind "the client made me do it".

The Measurement Problem

- The CFA GIPS standard says to treat taxes paid in the measurement period as an expense and just put a footnote regarding unrealized gain/losses left embedded in the portfolio.
  - The embedded gains/losses are ascribed no economic value. This is silly since there is a material probability that any embedded gains will require tax payments in future periods unless the investor is close to death.
- A related proposed standard encourages excess tax loss harvesting by assuming that if a portfolio has net capital losses in a tax year, the investor has realized capital gains in other investments
  - It has been estimated that only 8% of private investors are wealthy enough to make this assumption realistic. *Estimated tax alpha is upward biased.*
  - Warnings against overly aggressive tax loss harvesting were presented as early as 1998 by Paul Samuelson, Jr. Tax Sensitive Investing in the Long Run: Using a Simulation Environment to Quantify the Impacts of Different Investment Strategies (northinfo.com)

Full Liquidation and SIPE Methods

- Another proposed but equally flawed method is known as "full liquidation" accounting.
  - It is assumed that the portfolio is liquidated at the end of each performance measurement period to produce a "net value after taxes paid".
  - The after-tax return is computed as the percentage change in the "net values".
  - *This approach also produces upward biased returns* as the value of the tax payment isn't removed from the portfolio (i.e. put in escrow for future tax payments), so denominator used for return calculations is downward biased.

- In diBartolomeo (Journal of Performance Measurement, 2021) we propose an extension of the GIPS return calculation that includes an adjustment for changes in the present value of estimated future tax payments.
  - *This method realistically captures the maximum likelihood estimate of the after-tax return given uncertainty in future tax realizations and rates.*

The Tax Timing Option

- Another aspect of the problem that is ignored by the "tax budget" concept is *there is no consideration of the option value of properly timing the realization of tax losses.*
  - Whenever an investor could realize a tax loss, they also have the **option to wait** to possibly realize that loss at a later date.
  - Papers on the tax loss timing option include
    - Kalotay (JPM 2014), Kalotay (FAJ, 2016)
    - Kang, Paradise, and Dickson (FAJ, 2021)
    - Israelov and Lu (2022)
    - Davis, Li, and Nemtchinov (JBIS, 2022)

- As the value of the tax timing option is related to asset volatility, the interrelationship between portfolio risk and the "option to tax loss harvest later" can be built into the parameterization.

The Easy Solution For a Tax Budget Optimizations

- The simplest framework is to formulate the problem as a turnover constraint.
  - The procedure presented is approximate and as economically rational as we can make it.
  - We will call the turnover constraint for a specific optimization **T***
- The first step is to set a "budget" in tax dollars payable this year.
  - We will call the investor's chosen budget value **B**
- Next, optimize the portfolio with realistic tax rates, but on a pre-tax basis.
  - In the Northfield Optimizer, you can calculate taxes required but not impact the objective function by setting the "amortization constant" to zero.
  - We will call the amount of tax payments required to fully optimize on a pre-tax basis **P.**
  - We will call the amount of turnover required in the pre-tax optimization **Tp**

Calculating T* for an Optimization

- We can now parametrize T* (turnover constraint) as:

$$Z = (B/P) / M$$

$$T^* = Tp * MAX(1, Z)$$

M is the number of rebalancing events expected in the remainder of the tax year.

- Like all options, the tax timing option has time decay, so the solution is dependent on how much time is left in the tax year and the nature of the strategy (passive, active, value, momentum). Early in the year, the option to wait has more value so we want to slow down turnover.
- Given the gradient method used in the Northfield optimizer we are confident of maximum improvement in pre-tax utility.

Let's Do An Example

- We have $1 Million dollar account as of March 31, 2023.
  - We intend to rebalance monthly (2 down, 10 to go in the tax year) so **M is 10**.
  - The investor chooses a tax budget **B of $15,000**
- We optimize the portfolio pre-tax and determine that **P = $75,000** and **Tp = 30%**

$$Z = (\$15{,}000/\$75{,}000)/ (10) = .02$$

$$T^* = 30\% * .02 = .6\%$$

Let's Think Through What Happens April 30th?

- Since our optimization at March 31 was pre-tax, the expected value of taxes payable is **.6/30 * $75,000 = $1500, so the expected value of B at 4/30 is $13,500.**
  - However, we will know at what prices we traded at so we can exactly calculate taxes now payable for the trades in the 3/31 rebalance, so we can update how much is left of the $15000 budget to which we set **B.**
- Since prices will have changed between 3/31 and 4/30 (up, down, mixed), we rerun the pre-tax optimization to get updated values for **P and Tp**
  - We have nine optimizations scheduled for the rest of the year so **M = 9**
  - M will get progressively smaller which is consistent with decay in the tax timing option value, as allow more turnover per dollar of remaining budget.
- Calculate the new **T***
- Repeat the process for the rest of the year, B and M are going down.

An Exact Solution?

- It is also possible to structure an optimization to have an exact dollar amount of tax dollars payable or capital gains to be realized.
  - However, unless you're doing the last optimization of the tax year, the ability to solve for an exact amount is moot since price movements between subsequent rebalancing events are going to *wreck "the best laid plans of mice and men"*.

- While the details of this set up are beyond the scope of today's webinar, the set up is similar to the 2005 version of multi-account optimization which was described in our newsletter Northfield News- May 2005.pub (northinfo.com).

Conclusions

- We assert strongly that the rational objective function for private wealth investors includes tax effects as part of ex-ante risk adjusted returns.
- The idea of optimizing pre-tax subject to a "tax budget" constraint may be appealing to managers for its transparency to retail investors, but will produce consistently inferior results to a properly parameterized tax sensitive optimization.
- To the extent that managers are required by clients to operate within a tax budget irrespective of the flawed economics, *we have provided a procedure that frames the problem as iterative application of a turnover constraint through multiple rebalancing events.*
- Given the gradient solution used in the Northfield optimizer, users will achieve the highest possible improvement of the pre-tax utility function, conditional on the chosen turnover constraint.