Efficient Rebalancing of Taxable Portfolios

Sanjiv R. Das & Daniel Ostrov[1]
Santa Clara University

@JOIM
La Jolla, CA
April 2015

[1] Joint work with Dennis Yi Ding and Vincent Newell.

Problem Overview

An investor has a portfolio with a stock and cash position that can be traded periodically. The stock is subject to (very close to) the American taxation system. The portfolio has a given time horizon of T years.

- Basic question: What fraction, f, of the portfolio should be in stock?
- More specifically: What is the optimal static interval $[f^l, f^u]$ in which to dynamically maintain f over the portfolio's time horizon?
  - Maintaining f within an interval has been repeatedly shown to be the optimal strategy in cases with transaction costs, no taxes, and continuous trading. See, for example, Shreve and Soner (1994); Whaley and Wilmott (1997); Leland (2000); Atkinson and Mokkhavesa (2002); Janacek and Shreve (2004); Rogers (2004); Goodman and Ostrov (2010); and Dai, Liu, and Zhong (2011).
  - We will also consider an optimal dynamic interval where $[f^l, f^u]$ can change at time $\frac{T}{2}$.

Outline

1. Introduction
2. The Model
3. Base Case: Results and Observations
4. Alterations to the base case: Parameter Changes
5. Alterations to the base case: Model Changes
6. Conclusions

Previous approaches

1. PDE and Bellman Equation approaches: Require using the average cost basis for capital gains.
   - Dammon, Spatt, and Zhang (2001), Dammon, Spatt, and Zhang (2004), Gallmeyer, Kaniel, and Tompaidis (2006), and Tahar, Soner, and Touzi (2010).
2. Exact optimization with the full cost basis:
   - Dybvig and Koo (1996) - Maximum of 4 trading periods.
   - DeMiguel and Uppal (2005) - Maximum of 10 trading periods.
     - 1% certainty equivalent advantage over average tax basis.
   - Haugh, Iyengar, Wang (wp 2014) - Maximum of 20 assets with 20 periods using simulation and dual relaxation methods.

Our approach

New approach: Monte Carlo based optimization using 50,000 simulations of stock movement over T years.
Advantages:

1. Uses the full cost basis.
2. We used over 400 trading periods.
3. Can accommodate many more of the features of the American tax code than were previously possible.
4. Can accommodate any stochastic process for the stock movement.
5. Can easily extend to incorporate new features such as transaction costs or multiple stocks.
6. Yields many new economic insights, allowing us to answer (or more fully answer) many questions.

Some questions to be answered

Define

$$f^* = \frac{f^l + f^u}{2} \text{(optimal interval's midpoint)}$$
$$\Delta f = f^u - f^l \text{(optimal interval's width)}$$

1. Is f^* higher in a taxable or tax free account?
2. Is the "5/25" rule of thumb for rebalancing (i.e., $\Delta f = .10$) justified? Or is continually rebalancing better (i.e., $\Delta f = 0$)?
3. How sensitive are f^* and Δf to changes in underlying parameters?
    - Parameters: Stock and cash growth rates, Stock volatility, Investor risk aversion, Tax rates for capital losses and for capital gains, Portfolio size, Portfolio horizon T, and Trading period length.
    - E.g.: Should you increase or decrease f^* if the capital gains rate increases?
4. How do f^* and Δf change if the model changes?
    - Model changes: Introducing transaction costs, Investor being alive vs. deceased at time T, Allowing $[f^l, f^u]$ to change at time $T/2$, Using the average vs. full cost basis.
    - E.g.: How big of an advantage does using the full cost basis give over the average tax basis?

Assumptions and Notation: Assets

1. Two assets: risky stock and risk free cash.
   - Each can be bought and sold in any quantity, including non-integer amounts with, for the moment, negligible transaction costs.
   - Trading can only occur every h years. (E.g., $h = 0.25$ for quarterly trading.)
2. Stock: Stock evolves by geometric Brownian motion with a constant expected return, μ, and a constant volatility, σ. For simplicity, we do not consider dividends.
3. Cash: The tax-free (or post-tax) continuously compounded interest rate for the cash position, r, is assumed to be constant.

Assumptions and Notation: Taxes

1. Tax Rate: We assume the capital loss tax rate, τ_l, and the capital gains tax rate, τ_g, are constants that apply to both long term and short term gains/losses.
2. Wash Sales: We assume the presence of other stocks or stock indexes in our market with essentially the same value of μ and σ, so we can bypass wash sale rules. Given this, it is always optimal to immediately sell and rebuy any stock with losses. (Constantinides 1983 or Ostrov and Wong 2011).
3. Capital loss limits: No more than $3000 in net losses can be claimed at the end of each year. Net losses in excess of this amount are carried over to subsequent years.
4. Portfolio liquidation at time T: If the investor is alive, all capital gains are taxed. If the investor is deceased, all capital gains are forgiven and any remaining carried over capital losses are lost.

Optimization Goal and Trading Strategy

Optimization Goal: Determine the values for (1) the lower bound $f^l \geq 0$, (2) the upper bound $f^u \leq 1$, and (3) initial stock fraction $f^l \leq f^{init} \leq f^u$, that maximize the expected utility of the portfolio worth at a final liquidation time, T, using the following trading strategy at each trading period:

Optimization Goal and Trading Strategy

Optimization Goal: Determine the values for (1) the lower bound $f^l \geq 0$, (2) the upper bound $f^u \leq 1$, and (3) initial stock fraction $f^l \leq f^{init} \leq f^u$, that maximize the expected utility of the portfolio worth at a final liquidation time, T, using the following trading strategy at each trading period:

1. First: Sell and repurchase any stock with a loss. This generates money via a tax deduction.

Optimization Goal and Trading Strategy

Optimization Goal: Determine the values for (1) the lower bound $f^l \geq 0$, (2) the upper bound $f^u \leq 1$, and (3) initial stock fraction $f^l \leq f^{init} \leq f^u$, that maximize the expected utility of the portfolio worth at a final liquidation time, T, using the following trading strategy at each trading period:

1. First: Sell and repurchase any stock with a loss. This generates money via a tax deduction.
2. Second:
   - If the portfolio's stock fraction $f < f^l$, buy stock to raise fraction to f^l. Keep track of the cost basis for each purchase.

Optimization Goal and Trading Strategy

Optimization Goal: Determine the values for (1) the lower bound $f^l \geq 0$, (2) the upper bound $f^u \leq 1$, and (3) initial stock fraction $f^l \leq f^{init} \leq f^u$, that maximize the expected utility of the portfolio worth at a final liquidation time, T, using the following trading strategy at each trading period:

1. First: Sell and repurchase any stock with a loss. This generates money via a tax deduction.
2. Second:
   - If the portfolio's stock fraction $f < f^l$, buy stock to raise fraction to f^l. Keep track of the cost basis for each purchase.
   - If the portfolio's stock fraction $f > f^u$, sell stock to lower fraction to f^u. Sell stock with the highest basis to minimize immediate capital gains, which means LIFO in our model.

Utility

Power law utility function:

$$U(W_T) = \frac{(W_T)^{1-\alpha}}{1-\alpha},$$

where α is the coefficient of relative risk aversion and W_T is the portfolio's worth after liquidation at time T.

- In the case of no taxes like a Roth IRA: Optimal strategy (Merton (1992)) is $f = f_{\text{Merton}}$ where

$$f_{\text{Merton}} = \frac{\mu - r}{\alpha \cdot \sigma^2}.$$

That is, $f^l = f^{init} = f^u = f_{\text{Merton}}$ (or, equivalently, $f^* = f_{\text{Merton}}$ and $\Delta f = 0$). Note that f is constant in time and does not depend on the portfolio's worth.

Simulation Algorithm

1. For any given set of (f^l, f^u, f^{init}), the expected utility is approximated by the average utility over 50,000 Monte Carlo runs for the stock price evolution. This expected utility estimator is programmed in C, compiled and linked to be callable from the R programming language.
2. The optimization is run using a constrained optimizer in R, under Ubuntu Linux, for the three variables, f^l, f^u, and f^{init}, under the restriction: $0 \leq f^l \leq f^{init} \leq f^u \leq 1$.
3. We reran this using different sets of 50,000 simulations to check for consistency in the optimal values determined for f^l, f^u, and f^{init}.
4. The effect of varying f^{init} is quite small. Therefore, we only present the results for f^l and f^u. Generally, we will express f^l and f^u using

$$f^* = \frac{f^l + f^u}{2} \text{ and}$$
$$\Delta f = f^u - f^l.$$

Base Case Parameter Values

We will use the following "base case" parameter values:

- the stock growth rate, $\mu = 7\% = 0.07$ (per annum)
- the risk free rate, $r = 3\% = 0.03$ (per annum)
- the stock volatility, $\sigma = 20\% = 0.20$ (per annum)
- the risk aversion parameter, $\alpha = 1.5$ in our utility function
- the tax rate on losses, $\tau_l = 28\% = 0.28$
- the tax rate on gains, $\tau_g = 15\% = 0.15$
- the initial portfolio value, $W_0 = \$100,000$
- the time horizon before portfolio liquidation, $T = 40$ years
- quarterly rebalancing, $h = 0.25$ years

Later, we will experiment will alterations to each of these nine parameters.

Base Case Results: Roth vs. Taxable

For this base case:

$$\text{No tax (Roth IRA): } f^* = f_{\text{Merton}} = \frac{\mu - r}{\alpha \cdot \sigma^2} = \frac{2}{3} \qquad \Delta f = 0.$$

$$\text{Taxable account (investor alive at } T\text{): } \quad f^* = 0.71 \qquad \Delta f = 0.$$

Q1: Is f^* higher in a taxable or tax free account? Here, we see that it's higher in the taxable account. Why?

- *Not* because Roth cash is shielded from tax: The interest rate r is tax-free in both scenarios. If it weren't, f^* would be even higher than 0.71.
- It's because the capital loss rate, $\tau_l = 0.28$, is greater than the capital gain rate, $\tau_g = 0.15$!

Base Case Results: Investor lives vs. dies

If the investor dies when the portfolio is liquidated at time $T = 40$ years, all final capital gains are forgiven:

Taxable account (investor alive at T): $f^* = 0.71 \qquad \Delta f = 0.$

Taxable account (investor dies at T): $f^* = 0.76 \qquad \Delta f = 0.17.$

1. Forgiving capital gains makes stock more desirable, so f^* increases.
2. Increasing Δf reduces the number of transactions, which generally generates more capital gains at time T to be forgiven.

Even though living or dying at time T only affects the tax treatment at time T, it has a considerable effect on optimal long term investing strategy, especially on the optimal Δf.

Influences on the Optimal Interval Width Δf

Factors that affect the size of the optimal interval $\Delta f = f^u - f^l$:

1. Two factors that push Δf to be bigger:
   - The bigger Δf is, the more capital gains are deferred, which is advantageous even if the investor is alive and pays capital gains at the liquidation time T.
   - If the investor is deceased at T, then the bigger Δf is, the more gains are likely to be forgiven at time T.

Influences on the Optimal Interval Width Δf

Factors that affect the size of the optimal interval $\Delta f = f^u - f^l$:

1. Two factors that push Δf to be bigger:
    - The bigger Δf is, the more capital gains are deferred, which is advantageous even if the investor is alive and pays capital gains at the liquidation time T.
    - If the investor is deceased at T, then the bigger Δf is, the more gains are likely to be forgiven at time T.
2. Two factors that push Δf to be smaller:
    - The smaller Δf is, the closer we can keep the portfolio at or near the stock fraction that optimizes the expected utility.
    - The smaller Δf is, the more often we rebalance and the more likely we are to have losses, allowing us to take advantage of the fact that $\tau_l > \tau_g$ in current American tax law.

Effect of Parameter Changes

We next examine the sensitivity of the optimal stock fraction range, $[f^l, f^u]$, — or, equivalently, f^* and Δf — to varying the nine base case parameter values, one at a time.

We will work with the case where the investor dies at time T, since $\Delta f > 0$ in this case, permitting us to see the effect on Δf of changing parameters.

Varying the Stock Growth Rate from $\mu = 0.07$

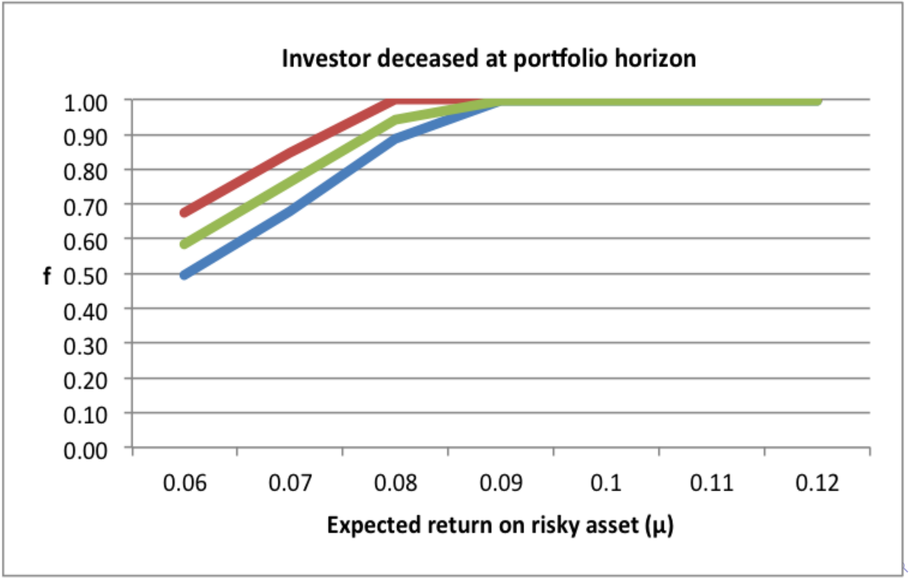

Varying the Cash Growth Rate from $r = 0.03$

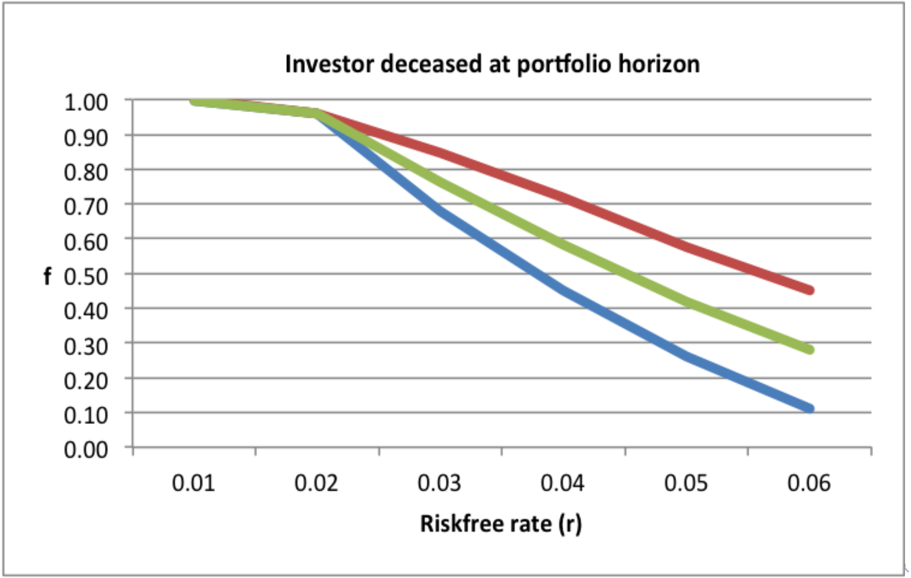

Constant risk premium $\mu - r = 0.04$

What if we increase both μ and r at the same rate?

Were there no taxes, there would be no change:

$$f_{\text{Merton}} = \frac{\mu - r}{\alpha \cdot \sigma^2} = \frac{2}{3}$$

What happens when there are taxes?

Varying μ, while holding $\mu - r = 0.04$

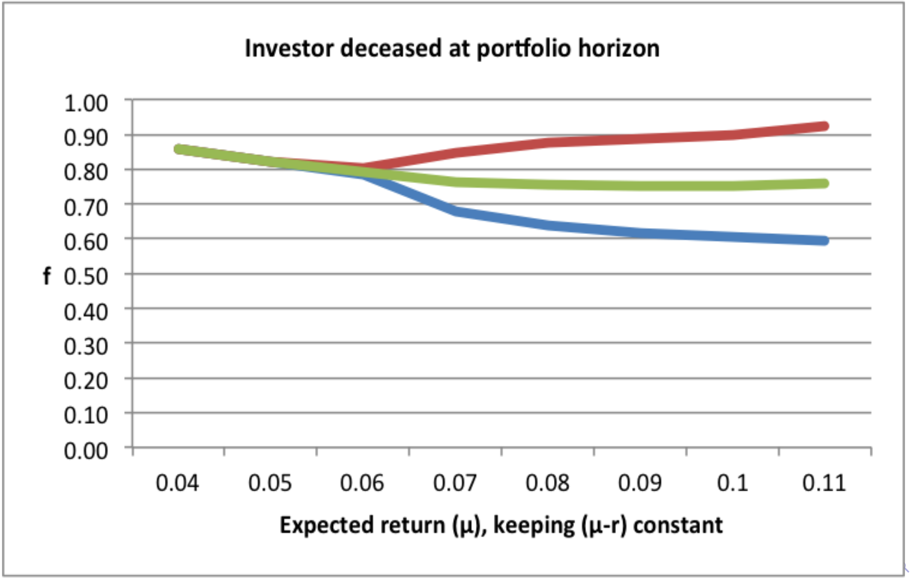

Comments on constant risk premia $\mu - r$

Why are there changes in the taxable case?

As μ increases there are more gains and fewer losses, which...

- ...reduces the ability to take advantage of the fact that $\tau_l > \tau_g$.
- ...makes it more advantageous to put off realizing gains.

Therefore, we see that

- Δf increases due to both of these effects.
- f^* decreases at first because of the decreased advantage of $\tau_l > \tau_g$, then increases as the growth in forgiven gains at death outweigh this factor.

Varying the Stock Volatility from $\sigma = 0.20$

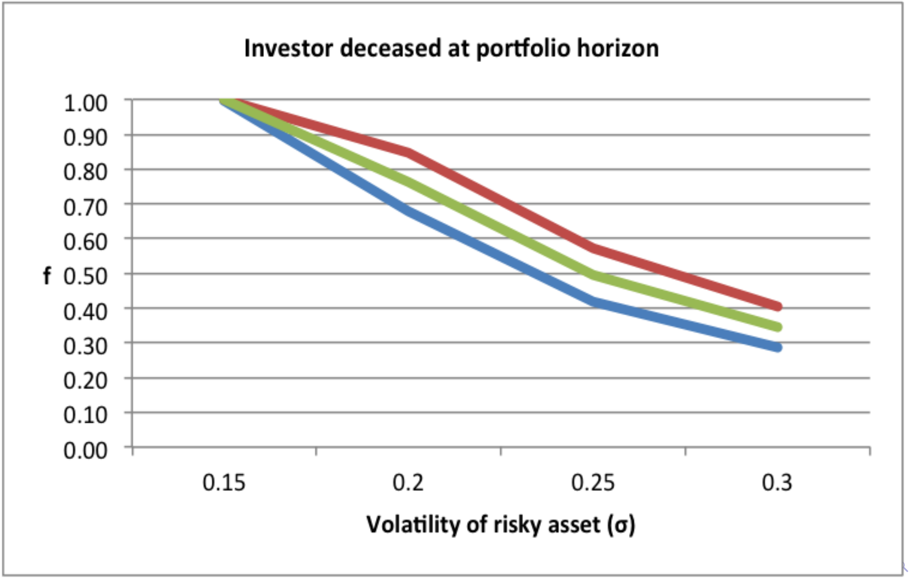

Varying the Risk Aversion from $\alpha = 1.5$

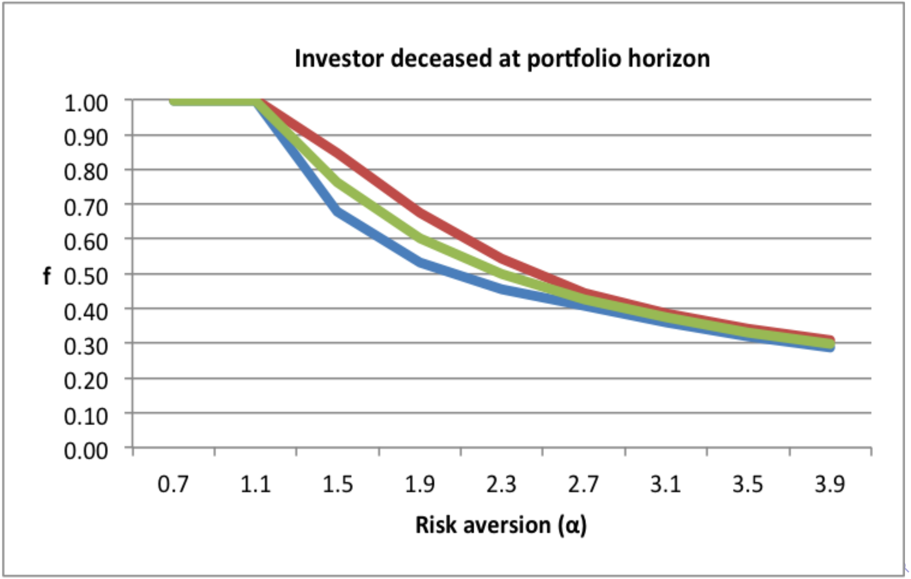

Comments on Risk and Aversion

1. No surprise that f^* decreases in both of these cases.
2. It is interesting that the effect of σ on Δf is so small.

Capital loss rate τ_l

If the capital loss rate, τ_l, increases, should you increase or decrease f^*?

Capital loss rate τ_l

If the capital loss rate, τ_l, increases, should you increase or decrease f^*?

Increase f^*: As the capital loss tax rate, τ_l, increases, stock becomes more desirable, so the stock fraction, f^*, should increase.

Capital loss rate τ_l

If the capital loss rate, τ_l, increases, should you increase or decrease f^*?

Increase f^*: As the capital loss tax rate, τ_l, increases, stock becomes more desirable, so the stock fraction, f^*, should increase.

Also decrease Δf to reset the cost basis more often, thereby generating more losses.

Varying the Capital Loss tax rate from $\tau_l = 0.28$

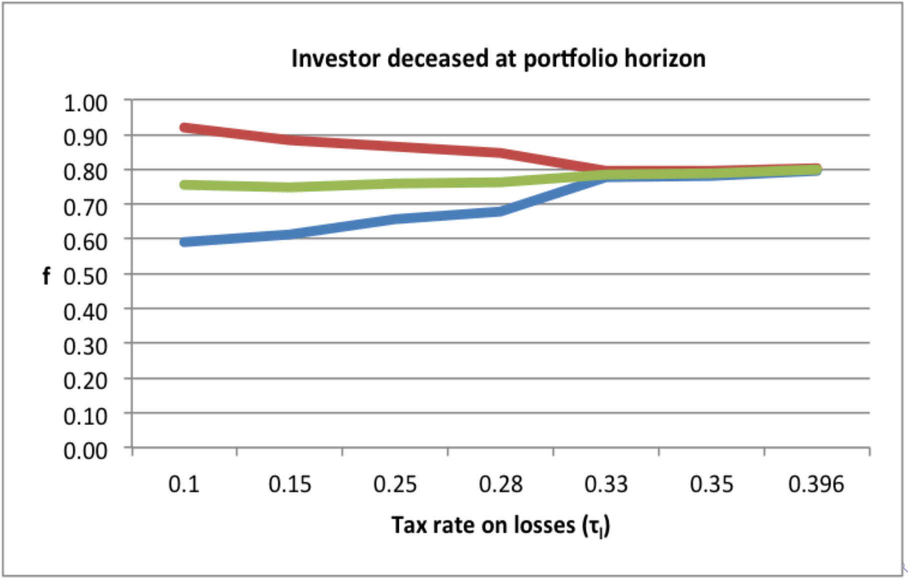

Capital gains rate τ_g

If the capital gains rate, τ_g, increases, should you increase or decrease f^*?

Capital gains rate τ_g

If the capital gains rate, τ_g, increases, should you increase or decrease f^*?

Decrease f^*: As the capital gains tax rate, τ_g, increases, stock becomes less desirable, so the stock fraction, f^*, should decrease.

Capital gains rate τ_g

If the capital gains rate, τ_g, increases, should you increase or decrease f^*?

Decrease f^*: As the capital gains tax rate, τ_g, increases, stock becomes less desirable, so the stock fraction, f^*, should decrease.

Right?

Capital gains rate τ_g

If the capital gains rate, τ_g, increases, should you increase or decrease f^*?

Decrease f^*: As the capital gains tax rate, τ_g, increases, stock becomes less desirable, so the stock fraction, f^*, should decrease.

Right?

Wrong!

Varying the Capital Gains tax rate from $\tau_g = 0.15$

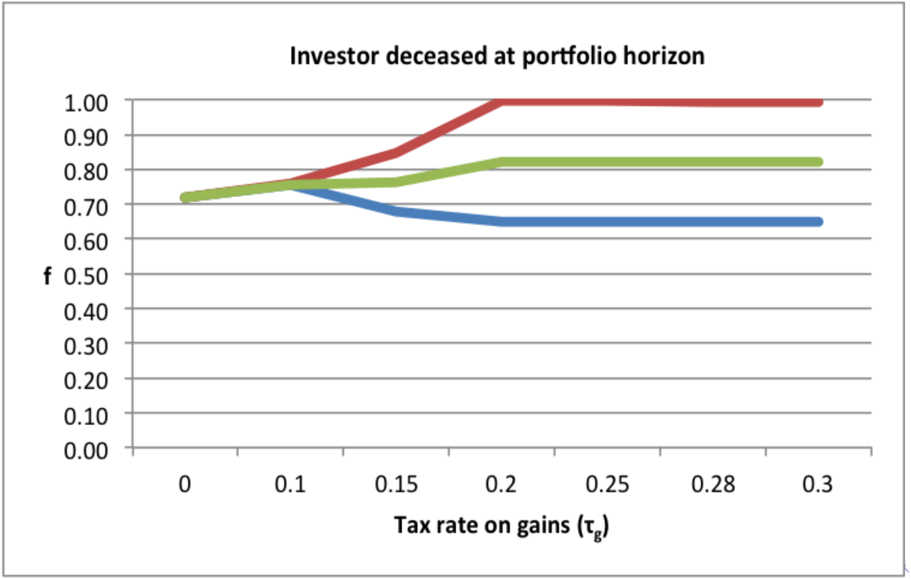

Comments on τ_g

Why?

1. Intuition: If $f = 1$, we have *no* capital gains. That is, *high* f means *low* capital gains in this extreme case.
2. More specifically: Suppose we have a portfolio with no stock volatility, a total worth of one dollar, and a stock fraction f. Our strategy is to annually rebalance the portfolio again to the stock fraction f. After a year, rebalancing generates $(\mu - r)f(1 - f)$ dollars of capital gains. This parabolic capital gains function equals 0 at $f = 0$, increases to its maximum value at $f = \frac{1}{2}$, and then decreases back to 0 at $f = 1$. So, if $f > \frac{1}{2}$, then *increasing* f *decreases* capital gains.

Not surprising: As τ_g increases, Δf increases to reduce realizing capital gains.

Initial Portfolio Worth, W_0

- For the case of no taxes:

$$f_{\text{Merton}} = \frac{\mu - r}{\alpha \cdot \sigma^2}.$$

  That is, there is no dependence on W_0.
- Were tax policy strictly dictated by proportional factors like τ_g and τ_l, the optimal strategy with taxes would be also be independent of W_0.
- However, the \$3000 limit on annual claimed losses is not a proportional factor.

Varying the Initial Portfolio Worth from $W_0 = \$100,000$

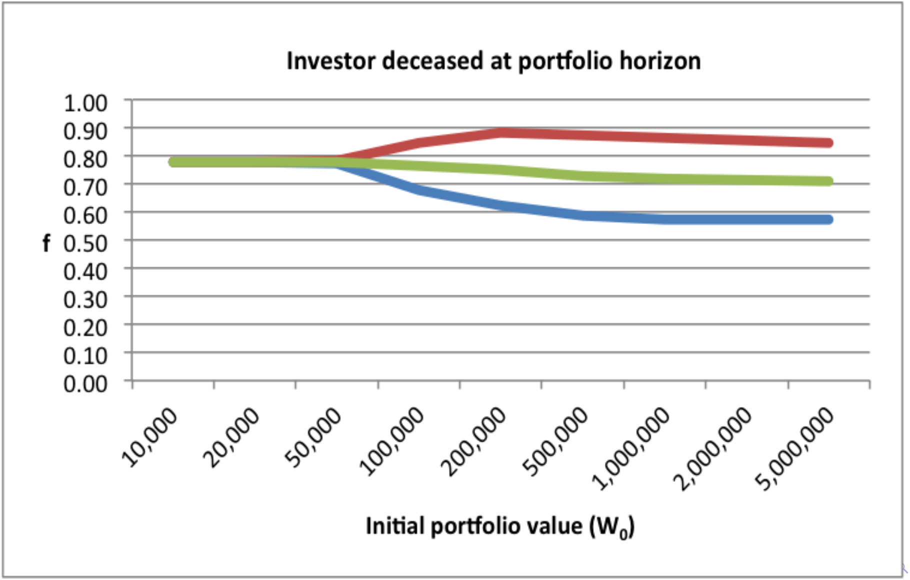

Comments on W_0

As W_0 increases, there is a mild decline in f^* due to the fact that the losses, as a proportion, become less useful as W_0 increases. Also, Δf increases, since creating losses becomes less useful.

Varying the Portfolio Horizon from $T = 40$ years

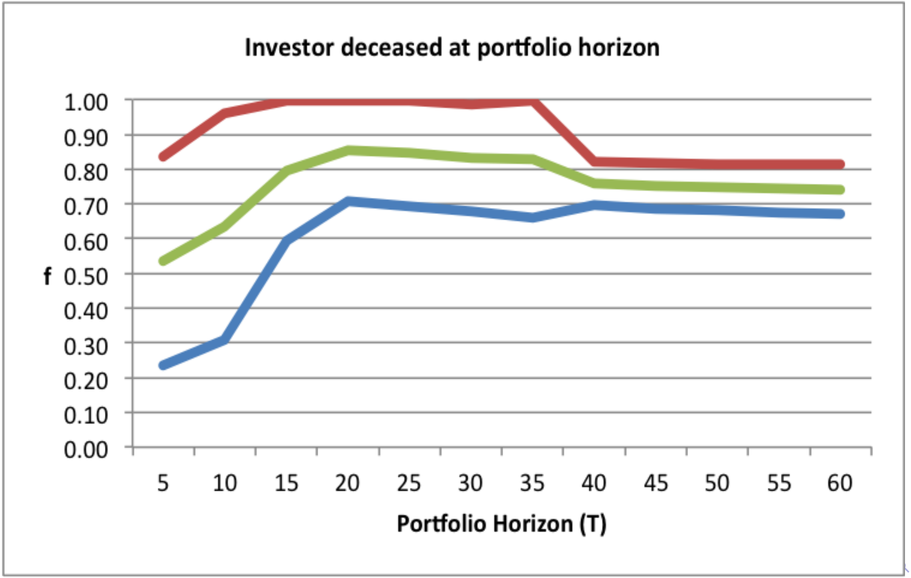

Comments on T

1. Initially, f^l and f^u (and therefore f^*) increase since the advantage of deferring/forgiving gains grows as T grows.
2. Then f^u decreases. Why?
   - If $f^u = 1$, we never sell stock with a gain. In this case, the bigger T is, the more likely f is to drift into the high end of the $[f^l, 1]$ interval. This portfolio has too much risk so reduce f^u.
   - If the stock price falls, the portfolio will do even worse, since stock is not sold when it reaches an $f^u < 1$, eliminating the tax loss option. This greater "wealth disparity" over longer time horizons is penalized by the concavity of the utility function, eventually forcing f^u to be reduced as T increases.

Varying the Rebalancing Period from $h = 0.25$ (quarterly)

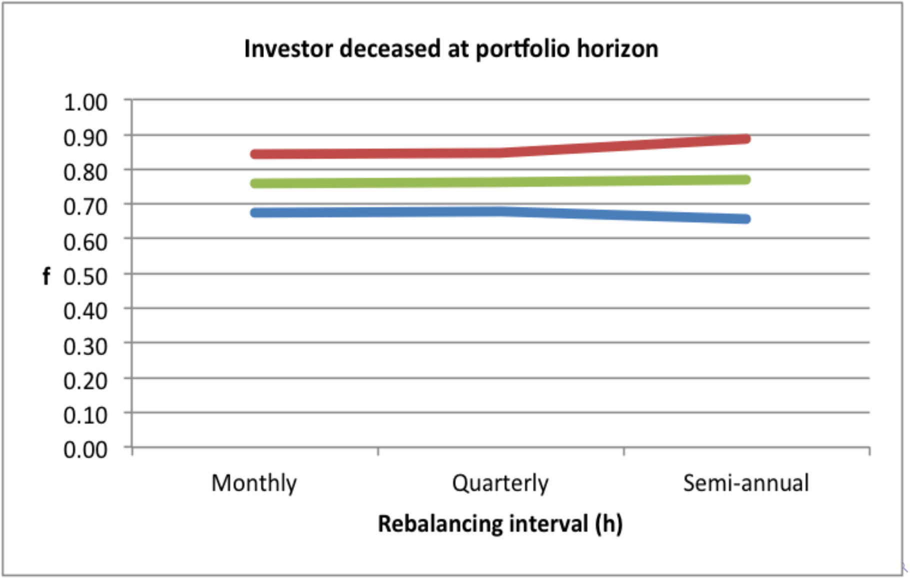

Comments on h

1. We see a slight increase in Δf as h grows, but the increase is small. The frequency of rebalancing is not a huge factor in the optimal strategy.

Changes to the Model

Our model is quite flexible. We next consider the effects of each of four different changes to our basic model:

1. We incorporate proportional transaction costs for buying or selling stock.
2. The investor is alive, instead of deceased, at the liquidation time T, so there are capital gains on the liquidated stock.
3. We allow the optimal stock fraction range, $[f^l, f^u]$, to change values when the portfolio is halfway to liquidation (i.e., at $\frac{T}{2} = 20$ years), instead of remaining constant.
4. We use the average cost basis instead of the full cost basis, which allows us to *quantitatively* measure of the suboptimality generated by the average cost basis.

Proportional Transaction Costs

Define e to be the proportion lost in transactions. That is,

$1 of cash buys $(1 − e) worth of stock.

$1 worth of stock can be sold for $(1 − e) of cash.

We still sell and repurchase any stock with a loss, even though this may no longer be optimal due to the transactions costs generated.

Expectation: As e increases, Δf will increase to reduce the number of transactions.

Varying the Transaction Costs from $e = 0$

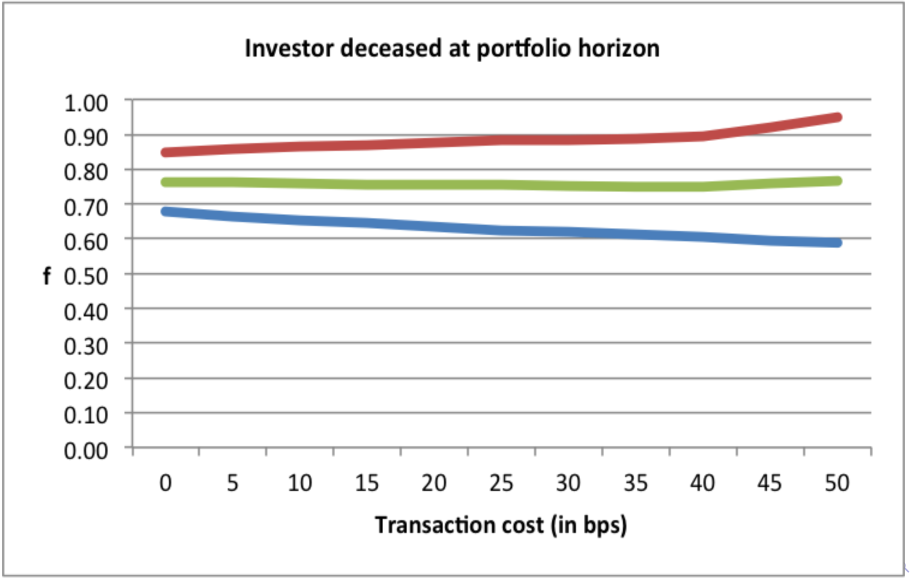

Effect of being alive vs. deceased at time T

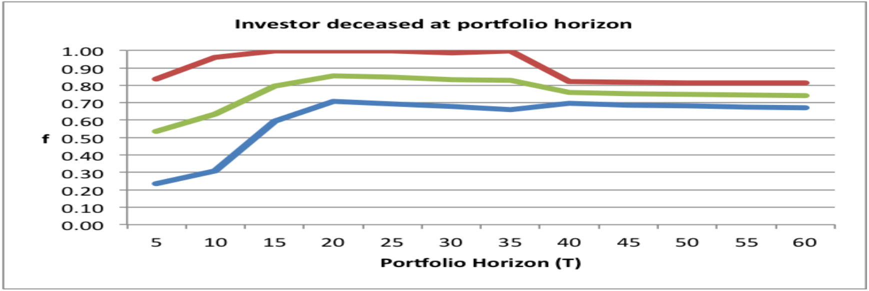

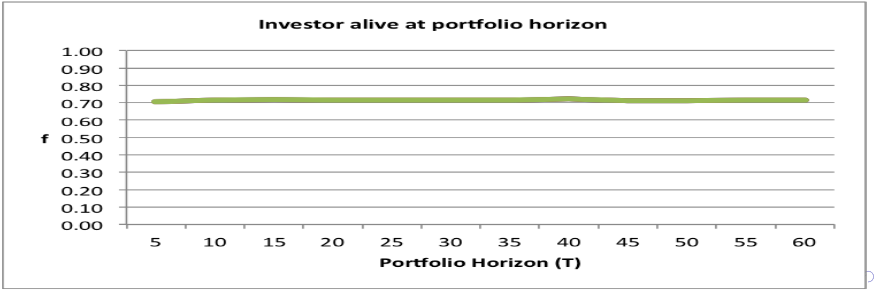

Comments on being alive at time T

Observations

- Even though being alive vs. deceased only affects the taxation rules at time T, the effects on the long term optimal strategy can be drastic!
- Note that these difference do not vanish as T grows.
- Only for the case of being alive: If $\Delta f = 0$ anywhere, as we see here, it is not surprising to see it remain zero for all T and for f^* to be constant for all T. You don't expect this necessarily in the case of being deceased.

Rebalance Region changes at $\frac{T}{2}$. (Dashed = Years 20–40)

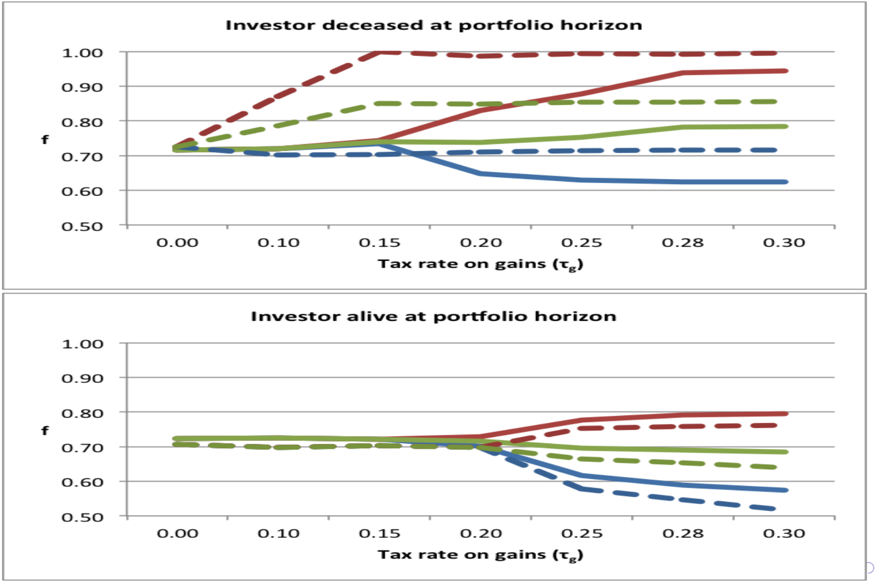

Comments on time-dependent strategy

- Optimize over five variables instead of three: f^{init} (the initial stock fraction), f^l and f^u for years 0–20, and f^l and f^u for years 20–40.
- When the investor is deceased, f^* increases in years 20–40 to take advantage of the forgiven gains at $T = 40$ years.
- When the investor is alive, f^* decreases a little in years 20–40, due to the inability to defer capital gains when they are forced to be realized during liquidation at $T = 40$ years.

Average cost basis

We simplified the model to use the average costs basis instead of the full cost basis.

We then compared the average cost basis results to the full cost basis results over each data point in each of the previous nine cases where we shifted a model parameter.

The results for shifting one these parameters, W_0, is on the next slide.

Average cost basis vs Full cost basis

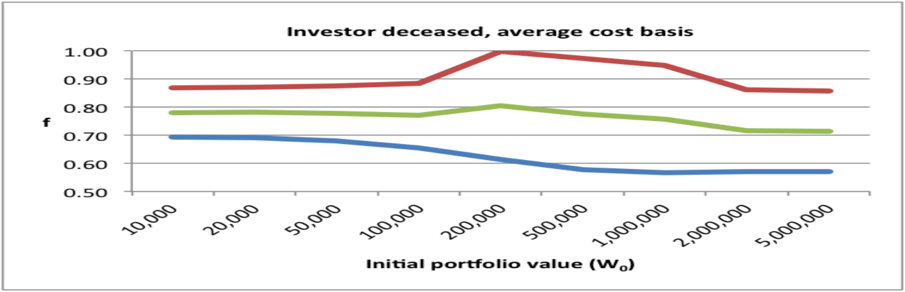

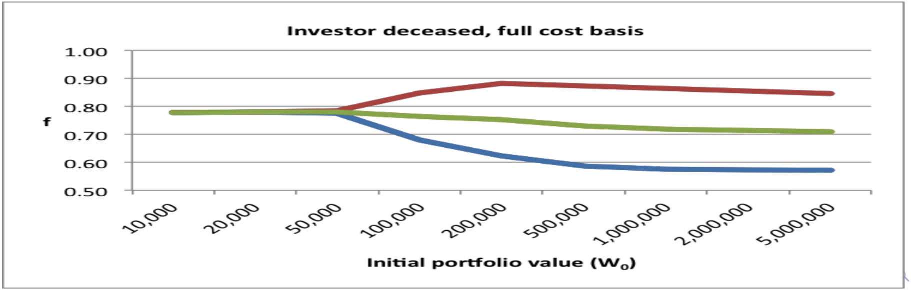

Comments on tax basis

Over all 9 parameter shifts:

- f^*: little change between average and full cost basis.
- Δf: shrinks in the full cost basis case by, on average, 6 percentage points. (Maximum in the data: 18.6 percentage points.) The shrunken Δf makes it more likely to generate losses from the high cost basis positions.
- But how much worse does the average cost basis case do compared to the full cost basis case?

Certainty Equivalent advantage

- Deceased vs. alive: Over the shifting nine parameters, being deceased has, on average, a 10% *certainty equivalent advantage* over being alive. That is, initially having \$1.00 invested with the capital gains forgiven at time T is as desirable to the investor as initially having \$1.10 and having to pay all capital gains taxes at time T. (The maximum certainty equivalent advantage in the data: 21.2%.)
- By comparison: Using the full cost basis has, on average, a 0.27% certainty equivalent advantage (with a maximum certainty equivalent advantage in the data of 1.12%). If the investor is alive, the average is 0.65% (with a maximum of 1.73%).
- Recall DeMiguel and Uppal (2005) found a 1% certainty equivalent advantage over average tax basis using 10 trading periods.

Results

Our Monte Carlo method yields a number of insights for investors using a taxable account:

1. Rebalancing rules such as the 5/25 rule are often not optimal, and should be used with caveats.
2. The frequency of rebalancing has little influence on the optimal strategy. Using a continuous time model vs. a discrete time model makes little difference.
3. Using the average cost basis vs. the full cost basis makes little difference. This justifies the use of the average cost basis needed in Bellman equation approaches.

Counter-Intuitive Results

A number of our conclusions are surprising:

1. The optimal f is often *higher* for taxable accounts than for tax-free accounts like the Roth IRA since $\tau_I > \tau_g$.
2. If the capital gains tax rate increases, then the fraction of the portfolio in stock, should be *raised*, not lowered.
3. Tax effects that only apply at liquidation can have a considerable effect on the optimal trading strategy, even though the strategy applies throughout the portfolio's lifetime from year 0 to year T.
   - This effect does not dissipate as $T \to \infty$.
   - But in contrast, the choice of f^{init}, the initial value of $f \in [f^l, f^u]$, has almost no effect on the performance of the portfolio, even when T is as small as 5 years.