January 2010

Transaction Costs and Equity Portfolio Capacity Analysis

Abstract

We describe the evolution of equity portfolio capacity analysis, and extend existing work to include the joint determination of turnover, return net of transaction costs, and capacity. Two new lessons emerge. First, previous focus on turnover as a choice variable is misplaced. Considering stock-specific transaction costs at the portfolio construction stage enables higher turnover levels, which themselves are determined through the interaction of alpha predictions and expected cost estimates. Second, excluding trading strategy from portfolio construction negatively impacts returns and reduces the potential capacity of the fund. The effective range of a fund's capacity in terms of assets under management is enlarged through explicit consideration of the implementation costs of an investment decision.

Yossi Brandes
Director
Investment Technology Group, Inc.

Ian Domowitz
Managing Director
Investment Technology Group, Inc.

Vitaly Serbin
Director
Investment Technology Group, Inc.

212.444.6300
info@itg.com
www.itg.com

We are grateful to Milan Borkovec and Lorelei Skillman for helpful comments, and to Michael Chigirinskiy for expert research assistance.

I. Introduction

In 1991, research by Perold and Salomon focused the attention of academicians and practitioners on the effect of the size of assets under management (AUM) on fund performance. The bulk of the ensuing literature documents an inverse relationship between a fund's size and its net return. The dominant explanation blames diminishing returns on rising turnover costs as fund size increases.

Turnover is required to exploit informational advantages, but performance deteriorates as fund managers turn over larger volumes of stock, incurring higher explicit and implicit transaction costs.[1] This observation leads to the *capacity problem*, or more generally, to *capacity analysis*, which has been defined in different ways over time. It may be understood simply to describe the dependence between equity portfolio performance and size of fund. Alternatively, capacity analysis may be the study of ways to increase a fund's assets under management for a given target level of return. The term has even been applied specifically to the identification of an optimal level of turnover for every level of assets under management. By any of these definitions, capacity analysis remains an important topic, used by plan sponsors in manager search, by chief investment officers to determine growth options, and by portfolio managers to calibrate turnover levels.

We build on the work of Kahn and Shaffer (2005) and Bull, Serbin and Zhu (2009), who are concerned with ways in which the range of a fund size might be extended while remaining efficient. We introduce the use of expected transaction costs, on a stock-specific basis, at the portfolio construction stage. In previous work, transaction costs are applied after portfolio optimization is completed, to simply compute net returns. Controlling for transaction costs while rebalancing the portfolio allows the maintenance of a significantly larger fund size, relative to previous practice, and permits a higher level of turnover than would be possible otherwise.

We begin with an overview of the current state of equity portfolio capacity analysis in Section II. As thinking advances from simply adding assets to a fixed portfolio strategy to consideration of stock-specific implementation costs, an explicit link is established among trading costs, portfolio optimization, and turnover. Consideration of implementation costs as an integral part of portfolio construction leads to the formulation of a Markowitz-style optimization problem in Section III.

The results of modeling net return at different levels of AUM are described in Section IV. Explicit controls for stock-specific expected transaction costs minimize adverse effects of fund size on net return, and increase the potential informational advantages accruing to higher turnover. The benefits of including expected costs in portfolio optimization increase with fund size itself.

Taxes and trading strategy are taken up in Section V. The role of taxes in portfolio construction is old ground, of course. We use this particular extension largely to reconcile differences in calculated turnover levels from our analysis with that observed in the mutual fund industry.[2] The role of trading strategy is more interesting. Accounting for expected implementation costs leads to turnover levels which strike a balance between the timely exploitation of new information and the avoidance of excessive costs. Trading strategy directly affects those costs, and therefore has a potentially important effect on capacity analysis. Although our basic results remain qualitatively the same for a wide class of expected cost estimates, quantitative conclusions can differ sharply with changes in trading

[1] See Wermers (2000) on informational advantages and Perold and Salomon (1991) and Vangelisti (2005) on performance.
[2] See, for example, Jeffrey and Arnott (1993).

strategy. We illustrate this point through a comparison of aggressive and passive trading behavior. Trading strategy becomes increasingly important as fund size grows.

Concept may be nice and theory even elegant, but putting ideas into practice often presents difficulties. Most issues in capacity analysis stem from a lack of complete data, even within a fund complex where one might otherwise expect such information to reside. We highlight some of the problems encountered in practical applications in Section VI. Pragmatic guidelines are offered with respect to the separation of trades resulting from simple cash flow considerations from those resulting from an investment decision.

II. Past is Prologue: Capacity Analysis Today

The current state of play

We have all seen announcements concerning the closure of funds to new investment, on the grounds that the limitation of further AUM growth preserves current fund performance. As academic authors and practitioners more fully documented the relationship between AUM growth and decline in realized return, attention shifted to explanations of underlying causes. The winner of the intellectual debate was the effect of implementation cost, and its close cousin, turnover, on fund performance. In hindsight, this was an obvious conclusion. After all, investment performance reflects two factors: the underlying investment strategy of the portfolio manager and the execution costs incurred in realizing those objectives.

The stock-specific nature of implementation costs led to a definition of fund capacity in terms of the capacity of an individual investment strategy. For fixed portfolio strategies, this point is illustrated in Figure 1.

Figure 1. Daily Market-Impact Costs for U.S. Indices Trade Lists up to $5 Billion

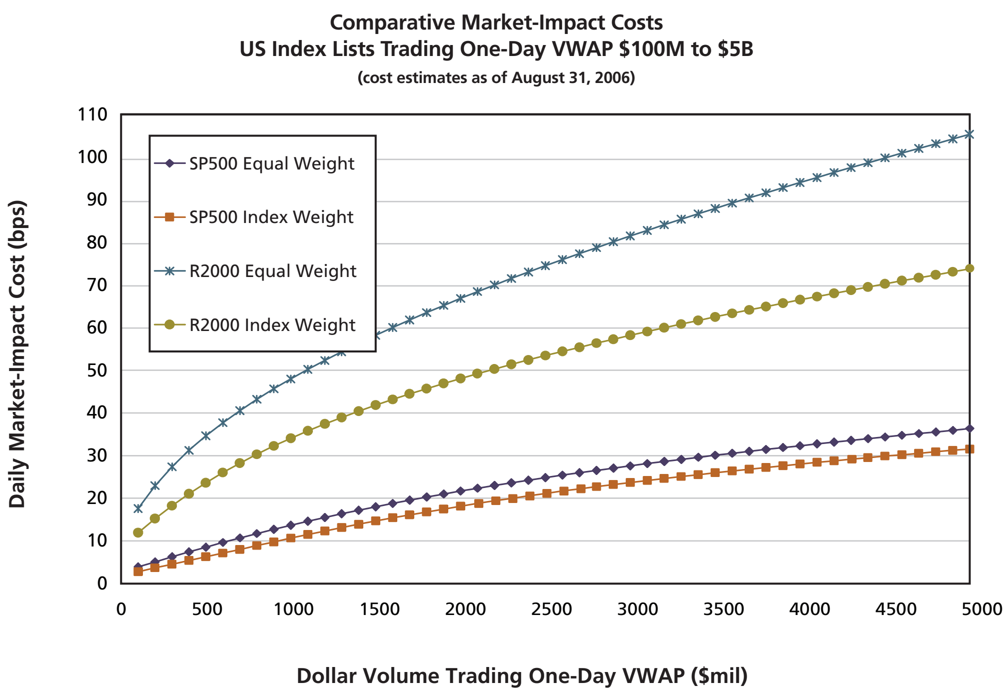

The figure is based on four trade lists representing distinct size-segments of the U.S. equity universe: equally-weighted and market-cap weighted large-cap stocks (SP500) as well as equally-weighted and market-cap weighted small-cap stocks (R2000). Trading an equally-weighted R2000 list costs noticeably more than an index-weighted list, as more dollar volume is pushed into the less liquid names. Asset classification is routinely influenced by this sort of evidence.[3]

Liquidity, measured in terms of implementation costs, can be affected by changing the breadth of the portfolio or by changing the investable universe. In turn, liquidity effects differ depending on the investment strategy, implying that capacity is a strategy-specific phenomenon.

Consideration of strategy has a logical extension to the method of implementation itself, namely trading. Given the evolution of transaction cost analysis (TCA), the incorporation of trading strategies into stock-specific trading cost estimates was slow. As a result, any focus on implementation cost in portfolio construction was on assumed fixed costs and turnover.[4] Even as such assumptions were replaced by better information over time, market-impact costs were taken into account only in the 'scoring' process to determine net returns, diminishing the importance of the trading strategy

[3] For example, the ordering of the international markets from least to most expensive in the fixed-dollar trade terms coincides with available floats. McDonald and Richardson (2006) suggest that index providers' decision to classify South Korea as an "emerging" market is largely for the convenience of institutional investors. With South Korea in their universe, emerging market funds can grow much larger without facing prohibitive turnover and maintenance costs, since South Korean stocks are much more liquid relative to stocks from other emerging markets.

[4] See Bogle (1994), for an early example.

proper. In other words, transaction costs were simply subtracted from expected portfolio returns after portfolio strategy in terms of names and quantities was determined. The emphasis remained on turnover, which was treated as the only real choice variable, conditional on an investment strategy and estimates of underlying implementation costs. Intuitively, this made sense, if only because turnover acted as a multiplier of costs and as an implicit penalty function in analytical frameworks.

One can get a long way with turnover. In Bull, Serbin and Zhu (2009), a long-short manager, using the S&P 500 as the investable universe, is modeled as an imperfect alpha forecaster, who rebalances monthly. In their framework, it is possible to derive optimal turnover levels for VWAP trading strategies, illustrated in Figure 2.

Figure 2. Optimal Turnover Rates for Different Levels of Assets Under Management (AUM)

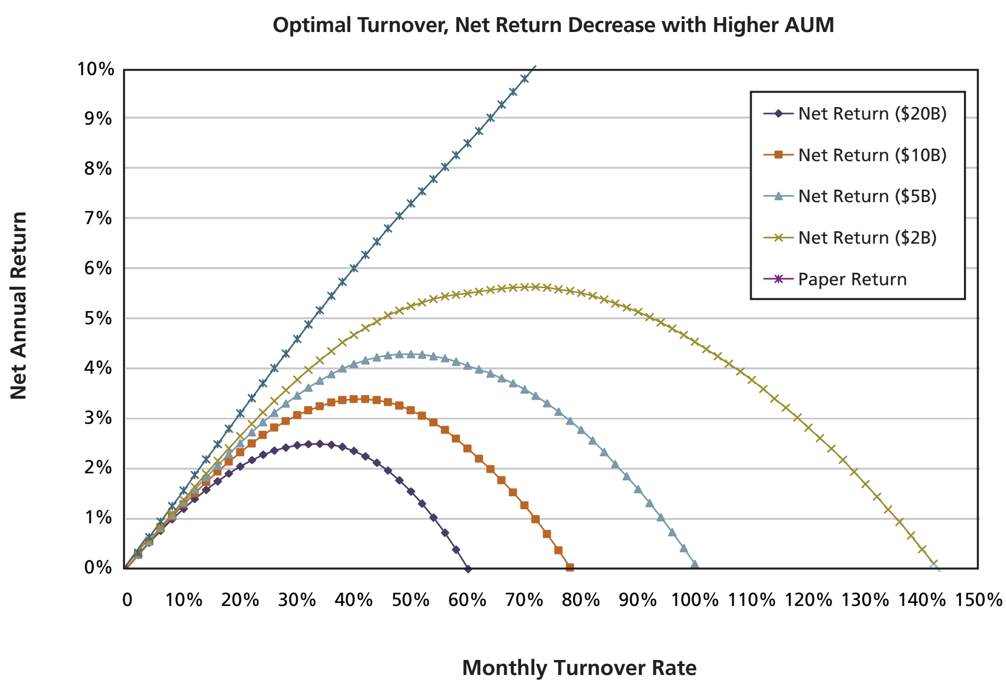

Due to transaction costs, the paper return dominates the returns of the portfolios at all AUM and turnover levels except zero. When turnover is very low, all four AUM levels exhibit almost a linear increase in net return. As turnover climbs past 20 percent, net returns for the larger AUM portfolios peak and drop quickly. For the $10 billion portfolio, net return is maximized at 3.40 percent with 40 percent turnover, while the $20 billion portfolio is maximized at 2.50 percent net return with 34 percent turnover. The $2 billion portfolio achieves a net return closer to 6 percent, and lower AUM levels approach the paper return, which includes neither commission costs nor market-impact costs.

This type of result motivates the idea of an efficient frontier for turnover, illustrated in Figure 3.

Figure 3. Turnover Optimality

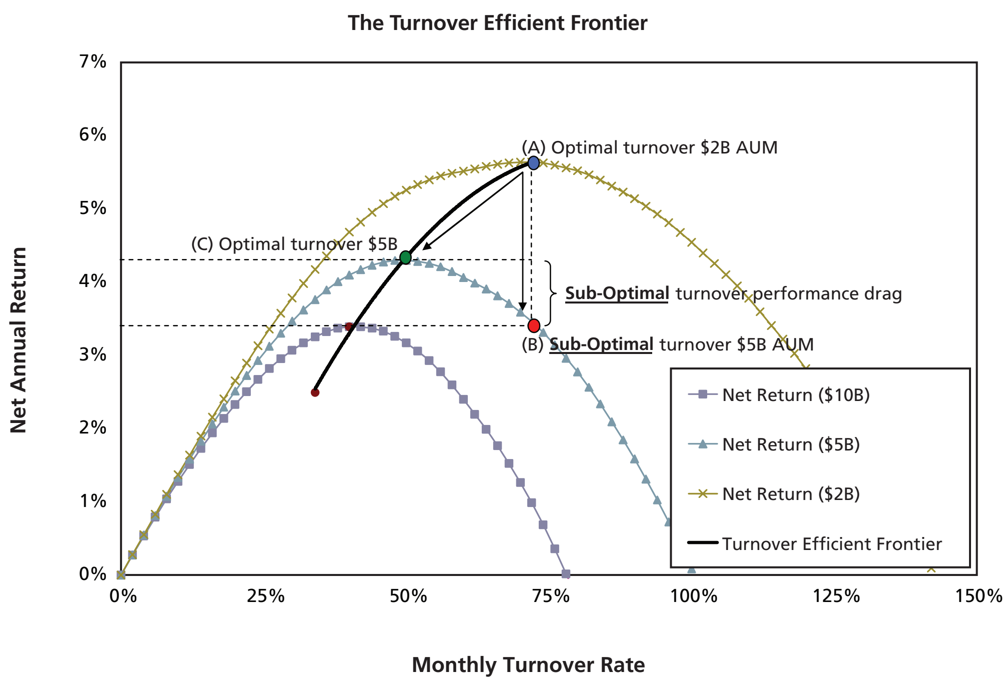

If turnover is held constant, the manager moves from point A straight down to point B. In order to achieve the highest return in this scenario, the manager decreases turnover to move to point C. Point B represents not just an inefficient portfolio but an inefficient *strategy* that is paying too much in market-impact costs from excessive turnover relative to skill level in forecasting alpha.

Isolation of causal factors in the form of transaction costs makes possible the idea that fund capacity analysis could be done ex ante. In short, expected implementation cost joins expected return in portfolio planning and execution.

The focus shifts as a consequence. The capacity problem can be recast as simple variant of the classical Markowitz optimization problem. The Fundamental Law of Active Management also is revisited by Coppejans and Madhavan (2007) and Bull, Serbin and Zhu (2009), from which some theoretical insights may be interpreted in capacity terms. For example, if the information coefficient is reinterpreted as the correlation between predicted and realized *net* alphas, then increasing AUM is accompanied by decreasing forecasting ability, and by a reduced correlation between risk-adjusted alphas and active portfolio weights.

At this point in the evolution of capacity analysis, turnover is a simple multiplier of the fixed costs of transacting. Bringing the turnover decision into the capacity problem requires some explicit link among trading costs, portfolio optimization, and turnover itself. We now turn briefly to the evidence that suggests such a link exists.

Portfolio Optimization and the Cost of Trading

The introduction of implementation costs into portfolio formation began with properly accounting for net, as opposed to "paper" returns after the composition of the portfolio was determined. In the simplest case, transaction cost analysis in portfolio construction is limited to the idea that costs simply eliminate part of the notional or "paper" return to an investment strategy, and therefore, costs should be controlled at the level of the trading desk only. Research on how to generalize and solve the portfolio construction problem in the presence of transaction costs dates all the way back to 1970, however.[5] The focus in most of the published papers has been on technical formulations and the mechanics of problem solving.

The economic implications of incorporating implementation costs directly into portfolio optimization are explored theoretically by Engle and Ferstenberg (2007) and on an empirical level by Borkovec, Domowitz, Kiernan and Serbin (2009). One of the optimization problems investigated in the latter work revolves around a two-year monthly rebalancing of a market and dollar neutral equity portfolio through August 2008, with the stock universe defined as the Russell 2000 Value Index. Stock-specific transaction costs are included in the formulation of the problem, forcing the various optimization stages to recognize implementation cost at each step. Details with respect to the precise problem, its solution, and results can be found in the reference. For our purposes here, Figures 4 and 5 yield links between the use of transaction costs as an explicit part of portfolio strategy construction and turnover. In the charts, we refer to the solution of the problem, including transaction costs ex ante, as being *cost aware*, while the *non-cost-aware* portfolio is optimized setting costs to zero and then subtracting expected costs from returns. The dotted lines delineate 95 percent statistical confidence bands.

[5] See, for instance, Pogue (1970).

Figure 4. Net Performance Based on Monthly Rebalancing

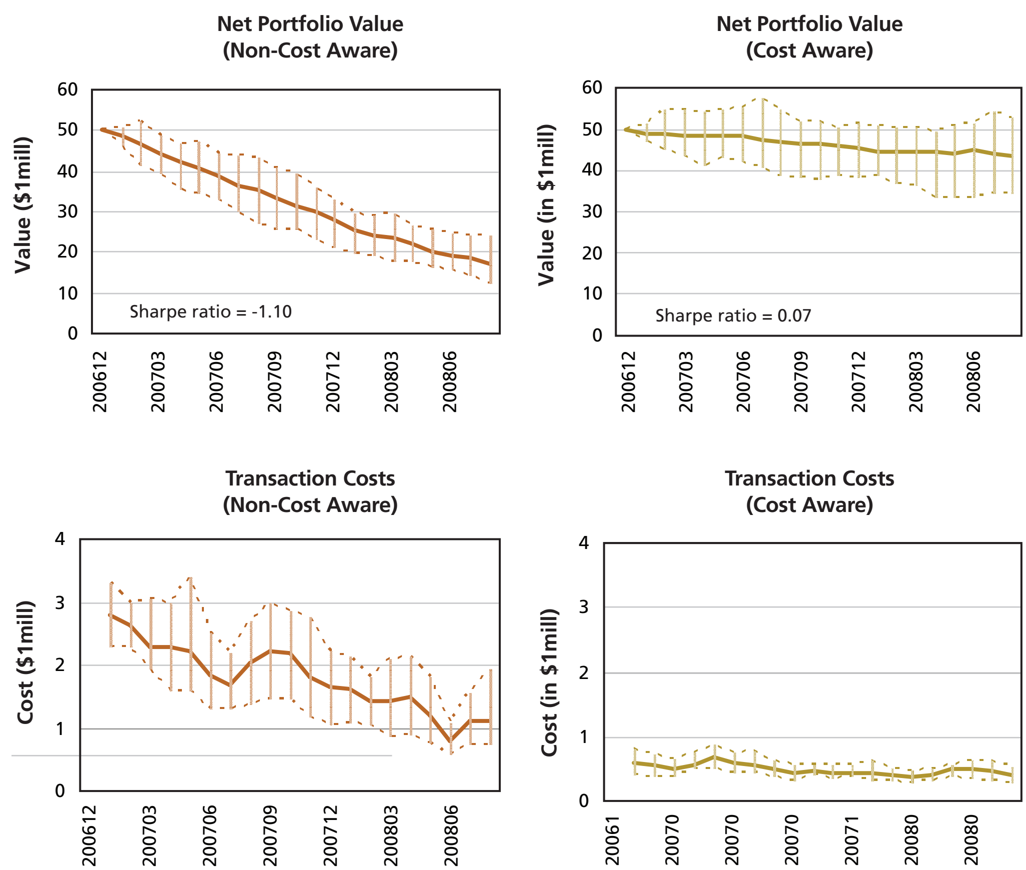

In Figure 4, the non-cost-aware rebalancing strategy loses on average more than 50 percent of its initial portfolio value by the end of the out-of-sample period, which translates into a negative Sharpe ratio of about -1.10. In contrast, the cost-aware portfolio performs much better, yielding a modest but positive Sharpe ratio of 0.07. The average dollar costs of implementing the cost-aware rebalancing strategy are two to four times lower than the costs of implementing the non-cost-aware strategy.

The costs of implementing the cost-aware rebalancing strategy remain lower despite the fact that its portfolio size starts to exceed the size of the non-cost-aware portfolio by as much as 50 percent in the last several months of the period considered. This is an important point in the context of capacity analysis. Another change shows up in turnover, depicted in Figure 5.

Figure 5. Turnover Comparisons Based on Monthly Rebalancing

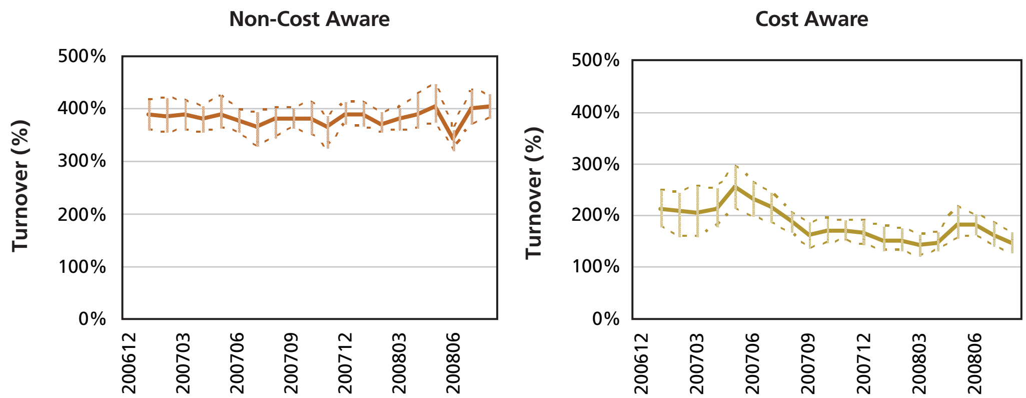

Turnover is calculated in percentage terms in this particular example, as dollars traded relative to the monthly portfolio basis. In a long-only portfolio, this statistic would have a maximum value of 200 percent. That value is doubled for the long-short portfolio with a 2:1 leverage ratio, i.e. the case in which the portfolio is 100 percent long and 100 percent short in dollar terms.[6] The total portfolio turnover goes down from almost 400 percent for the non-cost-aware to 250 percent and even 150 percent in the later months for the cost-aware rebalancing strategy. In other words, the cost-aware portfolio consistently exhibits significantly lower turnover and the monthly differences in this example can be as large as 270 percent.

We learn two lessons from this example, both applicable to capacity analysis. First, consideration of implementation costs as an integral part of portfolio construction matters, at least in the case of that amenable to a Markowitz-style methodology. In the context of capacity, similar integration offers the possibility of increasing fund size for any given level of required return. Second, the same method jointly determines expected cost, net returns, and turnover, and does so in such a way as to further reduce the cost to the portfolio. The implication here is what we have been looking for: turnover should not be considered a simple cost multiplier effect in a capacity analysis setting. We now turn to an investigation of both these points, beginning with the formulation of an appropriate optimization framework.

[6] In fact, the long-short portfolio turnover can occasionally exceed 400% when, for example, the short side outperforms the long side in the month following the optimization date, thus reducing the net portfolio basis.

III. The Optimization Problem

We consider a fund manager, who rebalances a portfolio every month subject to a typical set of constraints, augmented by penalties for turnover and trading transaction costs. Portfolio net returns are used as the yardstick by which we compare alternative combinations of turnover levels and fund sizes.

Every month, starting from December 2003 and ending in December 2008, we form a random portfolio of 100 stocks out of the eligible stock universe. The universe consists of all US stocks with a market capitalization exceeding $78 million, the approximate cutoff for Russell 2000 index as of the time of writing this article. We also require a trading price above $1 and no more than 10 missing returns in the last 60 months as of each portfolio formation date. This is admittedly an arbitrary choice, but any similar cutoff does not change the qualitative nature of the results. We add an additional restriction, that the half-spread for each stock does not exceed the 95th percentile of all half-spreads for the Russell 3000 universe for each month in the out-of-sample period. This requirement excludes extremely illiquid stocks, which are very expensive to trade. We consider six portfolio wealth levels ranging from $100 million to $5 billion, documented in Table 1.

Table 1. Parameter Values Used in the Optimization Problem

Portfolio wealth, W	**$100mln, $250mln, $500mln, $1bln, $2bln, $5bln**
Max. allocation weight, ϖ	**5%**
Monthly turnover bounds, t	**10%, 20%, 30%, 50%, 70%, 100%, 130%**
Annualized risk constraint σ	**35%**
Trading cost aversion coefficient τ	**0, 5, 10, 15, 20, 30, 50**

After the random portfolio is formed at the beginning of each month, we run the following optimization problem:

(1)

$$\max_{\omega} \; \hat{\mu}'\omega - \tau \frac{TC(W)}{W}$$

$$\omega'\Sigma\omega \le \sigma^2$$

$$\sum_{i=1,n} \omega_i = 1, 0 \le \omega \le \varpi$$

$$\sum_{i=1,n} \left| \omega_i - \omega_{i,0} \right| \le t$$

where $\omega_{i,0}$ is the initial allocation of the asset i, i=1,..100, ω^I is the vector of benchmark weights and the remaining notation is contained in Table 1 below.[7]

The vector μ in (1) represents the portfolio manager's estimate of expected returns. We simulate this alpha-forecasting process by drawing $\hat{\mu}$ from a statistical distribution given by:

(2)

$$\hat{\mu}_{i,t} = IC \cdot \mu_{i,t}^{observed} + \left[\Sigma_{\mu}^{1/2}\right]_{ii} \cdot \sqrt{1 - IC^2} \cdot \varepsilon_i, \quad i = 1, \ldots, n$$

where $\mu_{i,t} = R_{i,t} - r_f$, $R_{i,t}$ is the realized return for security i in month t; r_f is the risk-free rate; Σ_{μ} is the diagonal of the return covariance matrix; IC is the information coefficient (i.e. the proxy for the money manager's forecasting ability) and $\varepsilon \sim N(0,1)$. We report results for IC=5%, which roughly corresponds to an average forecasting ability.

Expected costs, TC, are derived from ITG's Agency Cost Estimator (ACE). Generation of these costs relies on a transparent methodology, and permits stock-specific estimates, which can be matched against expected returns.[8] ACE is a dynamic structural econometric model for the stock-specific expected price impact cost at the level of an individual order, differing by size of the order and the market conditions at the time of order submission. Permanent and transitory price impacts are explicitly modeled in such a way as to ensure that the first trade of a multi-trade order affects the prices of all subsequent sub-blocks sent to the market. These expected costs depend on the trading strategy. We assume a 10 percent volume participation rate strategy throughout. Using a different price impact model leads to similar qualitative conclusions, if the alternative model embodies a concave relationship between unit costs and volume. On the other hand, trading strategy affects implementation costs. Our results will suggest that taking expected costs into account matters a great deal, and trading strategy has a quantifiable effect on the investment strategy overall. We will, therefore, return to the role of trading strategy in Section V.

The covariance estimates are provided by the monthly U.S. model from the suite of ITG risk models. The model is estimated using time-series on a per-stock basis. The factor covariance matrix is scaled using an option-implied adjustment coefficient to provide forward-looking risk forecasts. The covariance matrix Σ of stock returns is computed as:

(3)

$$\Sigma = V^T F V + D$$

where F is the factor covariance matrix, V is the matrix of factor loadings, and D is the diagonal matrix of asset-specific variances.

Solving the optimization problem, we obtain an optimal portfolio which we hold for one month, record its return and then perform optimization again. At the end of December 2008, we have a time series of 60 monthly portfolio returns. We repeat this exercise 25 times. In other words, we draw 25 random portfolios and average the out-of-sample statistics across time (except for the first month when we start from cash) and across portfolios. We repeat this exercise for different combinations of parameter values presented in Table 1.

[7] We also considered a tracking-error constrained problem in which the first constraint in (1) is replaced by $(\omega_{i,0} - \omega^I)' \Sigma (\omega - \omega^I) \leq \vartheta^2$. These results are quite similar and not reported here.
[8] A complete description of the model, and results comparing expected to actual costs, are contained in ITG Inc. (2007).

IV. Trading Costs, Turnover, and Capacity

We model the interaction between turnover levels and net return at different levels of fund size. Since we introduce the transaction cost penalty T in the objective function, this relationship is described in terms of realized (as opposed to target) turnover.[9]

Figure 6 illustrates the central results of this exercise. Realized turnover is plotted along the horizontal axis and the associated out-of-sample net return appears along the vertical axis. We show this relationship for two levels of fund wealth ($250mln and $1Bn) and for three values of the transaction cost penalty coefficient (T=0, 20, and 50). The solid lines represent results corresponding to T=0, which in turn means that transaction cost estimates are applied to yield net returns only post-optimization, accounting for implementation cost but playing no direct role in portfolio construction. Dashed curves track results when such costs are introduced directly into the optimization problem, i.e., corresponding to T>0.

Figure 6. Return vs. Realized Turnover for Different Portfolio Wealth

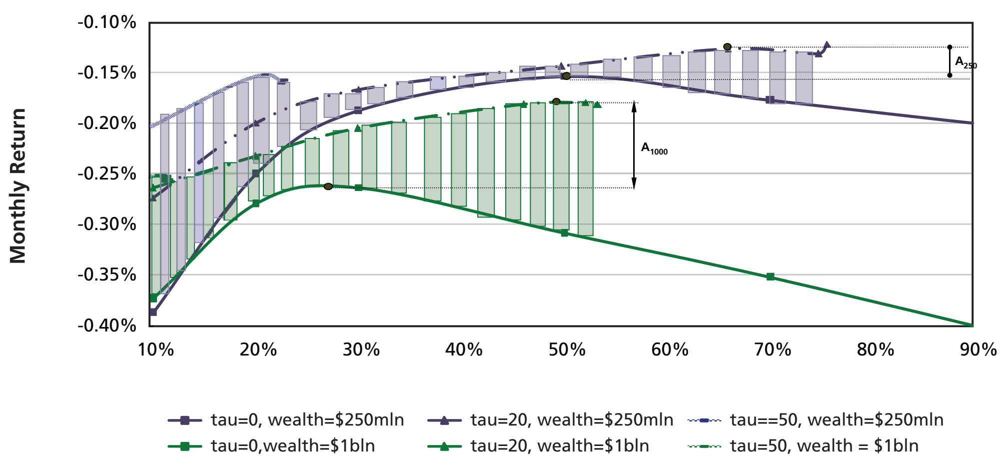

Increasing portfolio size leads to lower out-of-sample net returns. This result is not new, but the context is now important. Previous results implicitly contain the assumption that transaction costs enter only through the calculation of net return, post-optimization. Once that assumption is relaxed, the game changes markedly.

In particular, by picking an appropriate value for the cost aversion coefficient, τ, it is possible to minimize adverse affects of the fund size on net return. Applying transaction cost estimates post-optimization is inferior to including these costs when setting up the rebalancing

[9] Choosing high values for transaction cost penalty makes the turnover constraint non-binding. In that case, the composition of the optimal portfolio is determined mainly by the value of the transaction cost penalty, resulting in a virtually flat relationship between net returns and the turnover constraint.

problem. Comparing the solid blue and dashed green lines reveals that the net returns for a $1Bn fund with τ=20 are almost as high as the net returns for a $250mln fund which ignores transaction costs (τ=0).

To see this more clearly we mark the distance A, which is the difference between the highest net return with τ=0 and the highest net return with τ>0. This distance is the largest for the biggest portfolio (A_{1000} ~ 0.09% or almost 1.1% annualized) and the smallest for the smallest portfolio (A_{250} ~0.025%, over 0.3% annualized). More generally, the area between the solid and dotted lines in each color roughly represents the gain obtained by allowing τ to deviate from 0 (shown by the difference lines for $250mln and $1Bn wealth levels).

Comparing the size of the shaded areas, we conclude that the benefits of controlling for transaction cost increase with the fund size. Stock-specific transaction costs (and not turnover per se) effectively determine the capacity at which the fund remains profitable. The relative sizes of the shaded area also depend on the turnover range: for low turnover levels (up to 25%) the benefits of setting τ>0 are slightly larger for a smaller fund, while situation reverses when turnover levels exceed 25-30%.[10]

Transaction cost controls allow higher turnover levels, which translates into higher returns. For example, while optimal turnover levels for τ=0 are 18%, 14%, and 10% (for $250mln, $500mln, and $1Bn portfolios, respectively), these levels become 36%, 32%, and 25%, respectively, once τ is set to 20. Net return now peaks at a higher turnover level, which is consistent with the academic literature. That work identifies two major factors behind maintaining optimal turnover levels: the speed of release and homogeneity of information concerning the firm and transaction costs.[11] The positive relationship between the turnover level and net return can be related to information, as more trading is required to efficiently process new information. In practice this could be true only if the benefits of increased trading are not overwhelmed by accompanying transaction costs.

A higher net return is achieved not purely by increasing turnover, but by doing so in a cost-effective way. This is illustrated in Figure 7, which shows the median trading volume associated with portfolio rebalancing (expressed as a percentage of average daily dollar volume, ADDV) versus realized turnover. The slope of the line decreases dramatically when the coefficient of cost aversion is increased. While the slope is close to 1 for τ=0, it is virtually zero when τ=30, measuring the slope over the turnover range between 0 and 30% per month. Achieving this type of result might include, for example, selection of more liquid stocks with similar alpha characteristics and spreading the portfolio across more names. This helps to achieve the benefits of running a high-turnover strategy while avoiding to a large extent transaction costs associated with it. Our findings echo the theoretical predictions of Coppejans and Madhavan (2007), who note that controlling for transaction cost helps to maintain portfolio breadth and turnover, both of which have a positive impact on alpha.

[10] Actual turnover constraints used to plot Figure 1 are different from realized turnover values shown on the chart. The 1-to-1 correspondence between realized and actual turnover holds only for τ=0 line. While the same number of turnover constraints was used for τ>0, the resulting lines span shorter range of realized turnover than τ=0 line.

[11] See, for example, Karpoff (1986) and Domowitz, Glen & Madhavan (2001).

Figure 7. Median Trading Volume, % of ADDV vs. Realized Turnover, $1Bn Portfolio

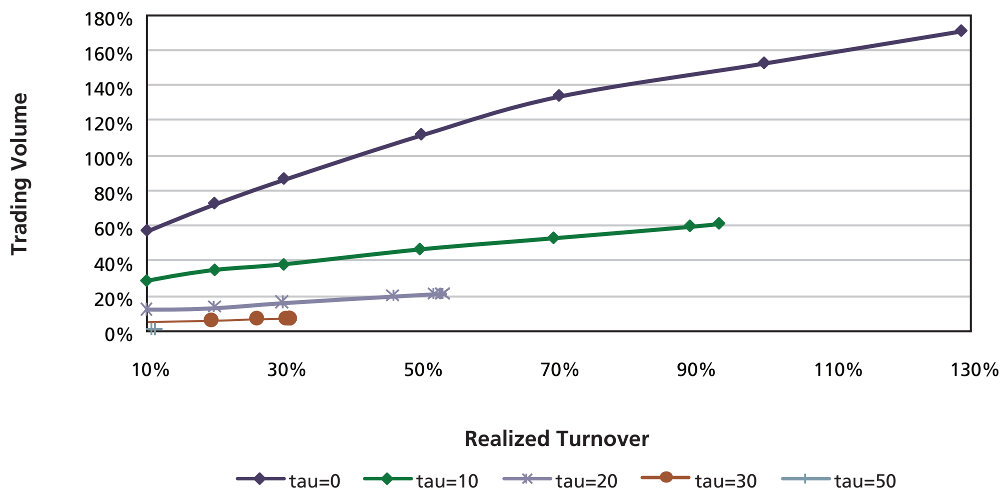

Higher turnover levels are not always possible, however. Some of the curves in Figures 6 and 7 are shorter than the others. The shortest ones correspond to the highest values of τ. For instance, with τ=20 a $250mln portfolio can increase its turnover to 75%, while a $1Bn portfolio with τ=20 can grow turnover only to 55%. In other words, setting τ too high limits the turnover range as a large coefficient of cost aversion completely dominates turnover constraints. With high values of τ (e.g. τ=50), realized turnover can grow to only around 20% per month for a $250mln portfolio and to only 10% per month for a $1Bn portfolio.

It is instructive to look at how turnover varies with the wealth level. For the purpose of this discussion, optimal turnover is defined as the level corresponding to the highest out-of-sample net return. Figure 8 depicts this definition of optimal turnover as a function of fund size. Three turnover levels are shown: the green line corresponds to the optimal coefficient of cost aversion, the blue line corresponds to the absence of cost control (coefficient of cost aversion T is set to zero) and the red line corresponds to τ=50. We also show the corresponding net returns. Two observations are immediately visible from the chart.

Figure 8. Optimal Turnover for a Given tau and Portfolio Wealth

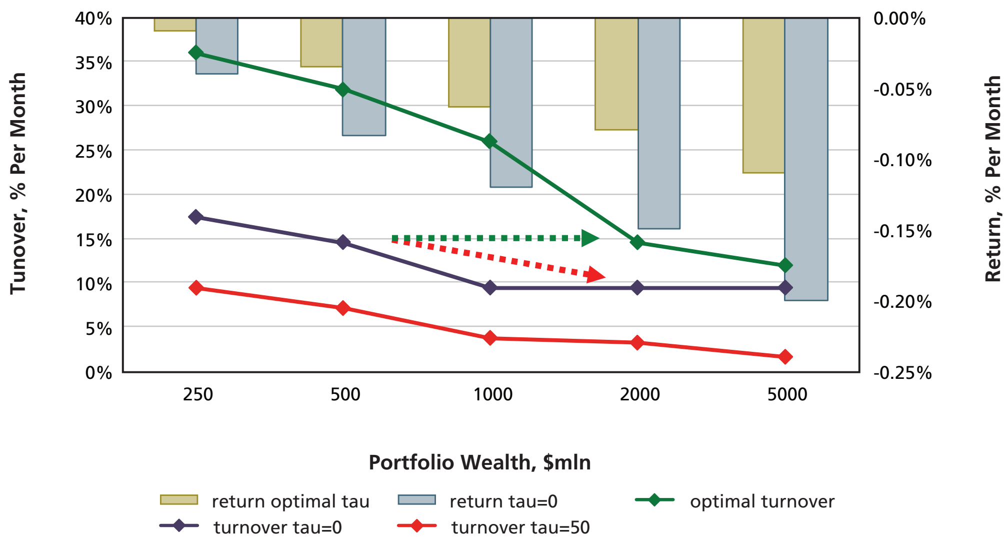

First, the optimal turnover level goes down as the fund size increases. This is intuitive and reflects other results available in the literature. In the broader framework considered here, such a result must be interpreted only in the context of a fixed level of implementation cost penalty, however. The implications of varying that penalty bring us to the second point.

Higher turnover levels are achieved by picking the right transaction cost penalty at each particular level of the fund wealth. For example, suppose that the fund manager does not include transaction costs into the rebalancing problem for a $500mln fund. In order to increase the fund size to $2Bn and still maximize net return, the manager should decrease turnover from 15% per month to 10% while giving up some net return (from -0.04% to -0.08% per month). We depict this move by the dashed red line. However, if growing the fund were accompanied by introducing effective cost controls (setting τ=20), the fund manager could keep turnover at 15% per month, which preserves net return. This transition is depicted by the dashed green line. Careful modeling of transaction costs during portfolio rebalancing allows reaping the informational benefits of maintaining the same turnover level with more wealth, extending the effective range of the fund's capacity.

Figure 9 provides another illustration that the effect of turnover on the fund's net returns and capacity follows from the more fundamental link originating in transaction cost control. There, we show the envelope curve which traverses all possible turnover levels for a $500mln fund and which passes through the highest net return for each of them. The different colors of the curve correspond to the value of the coefficient of cost aversion which maximizes the net return for a particular range of turnover.

Figure 9. Pre Tax Net Return vs. Realized Turnover

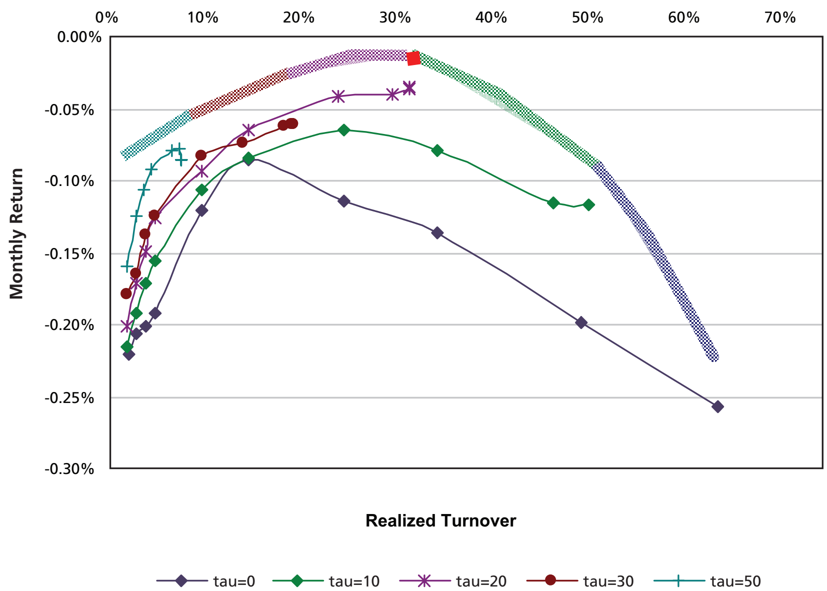

Maximization of net return solely through turnover is not generally possible. Without adjusting the coefficient of cost aversion, there is no guarantee that any chosen level of turnover would result in the highest net return. Fixing the cost aversion coefficient and then shifting the turnover level up or down also would not work, because the optimal cost aversion coefficient is different for each turnover level. For example, when the turnover level is 10-20% the optimal τ is 50, but it becomes 30 for a turnover range between 20% and 45%. A global optimum in terms of portfolio performance is achieved by setting τ=20, with the maximum supported turnover of about 30-32%. We can easily reconcile this point (marked in red) with the corresponding point on Figure 8, representing the optimal turnover level for a $500mn portfolio. Running the combination of 30-32% turnover and τ=20 strikes a balance between utilizing informational advantages that come with a higher turnover level while guarding against excessive implementation costs. While the envelope in Figure 9 is for a $500mln portfolio, the curves and the argument for different portfolio wealth levels are similar.

In order to achieve the highest possible return for a given portfolio and fund wealth, the manager may ignore turnover constraints, and improve performance, if costs are managed through the coefficient of cost aversion. Figure 10 depicts this idea for $500mln fund. Traversing through the values of cost aversion coefficient ensures traversing through all relevant turnover levels. Only return corresponding to the highest possible realized turnover level for each value of T (where we would end up in the absence of turnover constraint) is shown. The values of T are marked on the chart.[12] Varying only the cost penalty coefficient we are able to arrive at the same sweet spot for a $500mln fund (which is T=20, and turnover ~32%) as we did by varying *both* the cost penalty and turnover

[12] For convenience we mark in red the points on Chart 9 and Chart 10 that correspond to each other (for T=20).

constraint. Since turnover levels for the optimization problem illustrated on Chart 10 were never set explicitly, this illustrates the point that turnover is determined jointly with net returns. Expansion of fund capacity should advance through the direct control of stock-specific transaction costs.

Figure 10. Monthly Return vs. Realized Turnover

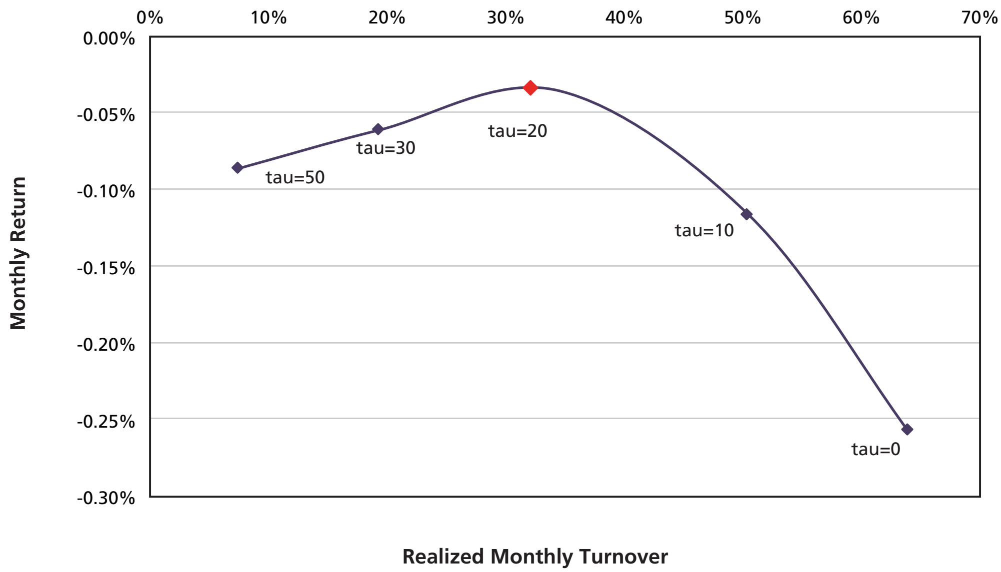

V. Taxes, Trading Strategy, and Turnover

A Role for Taxes

Optimal turnover levels reported in the previous section are higher than the typical turnover levels observed in the mutual fund industry. Average fund sizes in the large-cap blend, mid-cap blend, and small-cap blend categories are $2.3Bn, $850mln and $600mln, respectively and the average turnover levels are 80%, 90%, and 95% per year, respectively.[13] Simulation runs for a $500mln fund reported in the previous section indicate an optimal turnover level of 30 to 32% per month. There are features of our 'experimental design' which may account for this, of course, including but not limited to over-optimism with respect to innovations in the alpha generating process on a month-by-month basis.

We offer another reconciliation of the difference in turnover levels by considering taxes. This point has been made by Jeffrey and Arnott (1993), who demonstrate that realized capital gains taxes play a substantial role in the calculation of net returns.[14]

[13] The source is Morningstar, using information taken from their web site, www.morningstar.com.

[14] They find, for example, that increasing turnover from 25 percent to 50 percent requires an additional 63 basis points in return to offset the increase in taxes, and an additional 45 basis points in return to offset the increase in taxes when moving from 50 percent to 100 percent turnover.

Assuming a combined federal, state, and local tax rate of 35% and a 6% annual asset growth, those authors estimate that an investor would lose 2.1% per year to taxes if the fund's turnover is kept at 100%. Going back to the basic problem setup described in Section III, we change the way we compute realized out-of-sample net returns by approximating taxes to be paid on realized capital gains. We do not consider taxes at the optimization stage; we simply subtract taxes when calculating out-of-sample return. We will return briefly to this point below.

We assume capital gain taxes to be a simple function of turnover. We calculate turnover as it is defined by the SEC, i.e. the lesser of purchases or sales for one year divided by the average of portfolio wealth during this year.[15] Denote total portfolio purchases in dollars in year y as:

(4a)

$$b^y = \sum_{i=1,..,n} \sum_{m=1,..,12} b^y_{i,m}$$

and total portfolio sales in dollars as:

(4b)

$$s^y = \sum_{i=1,..,n} \sum_{m=1,..,12} s^y_{i,m}$$

where $b^y_{i,m}$ and $s^y_{i,m}$ are the dollar amounts of asset i bought and sold respectively within month m of year y.

We then aggregate monthly portfolio returns r_t in order to arrive at the fund wealth at the end of year y (or at the beginning of year y+1):

(5)

$$c^{y+1} = c^y \prod_{t=1}^{12} (1 + r_t)$$

where C_y and C_{y+1} the portfolio wealth at the beginning of the year y and y+1 respectively.

Taking (4)-(5) into consideration, we calculate turnover in year y as:

(6)

$$to^y = \min\left(\frac{b^y}{(C^y + C^{y+1})/2}, \frac{s^y}{(C^y + C^{y+1})/2} \right)$$

[15] Since the goal of this section is to compare simulated turnover levels with the levels observed in the mutual fund industry, we now adhere to a slightly simplified version of the SEC definition on the Form N-3 which is required to be filed by investment management companies under the Securities Act of 1933 and the Investment Company Act of 1940. It is different from the values shown in the earlier sections where we calculated turnover per the last constraint in (1).

While capital gain taxes depend on the holding period, we assume that the taxable part of realized capital gain for year y is proportional to the annual turnover in that year. This assumption is consistent with the one used in Jeffrey and Arnott (1993) in which the authors assume a "hockey stick" relationship between turnover and holding period. At the same time, the taxable part of the capital gain cannot be more than 100% of the total. We approximate taxes G^y for year y as:

(7)

$$G^y = g^y (C^{y+1} - C^y), \text{ if } (C^{y+1} - C^y) > 0$$

$$G^y = 0, \text{ if } (C^{y+1} - C^y), \leq 0$$

where $g^y = to^y$ if $to^y < 1$ and $g^y = 1$ if $to^y \geq 1$.

Using a tax rate of 15% (t_{rate}=0.15), we calculate the realized annual net return as:

(8)

$$R^y = \left(\frac{C^{y+1} - C^y - t_{rate} G^y}{C^y} \right)$$

At the end of December 2008, we have a time series of five annual portfolio returns. Repeating this exercise 25 times, we have 25 random portfolios whose out-of-sample statistics we average across time (except for the first month when we start from cash) and across portfolios. We repeat this exercise for different combinations of parameter values presented in Table 1.

We present only a simple illustration of how taxes affect the results. In Figure 11 we show pre- and post-tax returns for two combinations of fund size and cost aversion coefficient: \$500mln/T=50 and \$5Bn/T=30.[16]

[16] The evidence for other fund sizes and other parameter values is qualitatively similar and not reported here.

Figure 11. Net Return vs. Realized Turnover

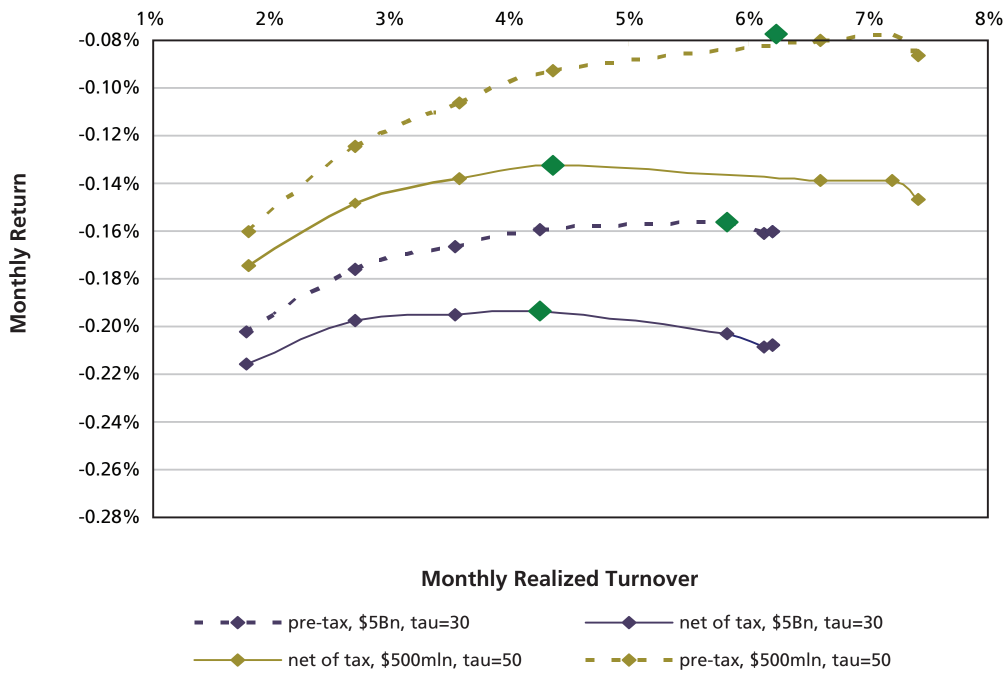

Applying taxes reduces net return by design, of course. It also reduces the optimal level of turnover: from 5.8% per month to 4.25% per month for a $5Bn fund and from 7.2% to 4% for a $500mln fund. Both values of after-tax turnover are now in line with the turnover prevailing in the industry.

Taking taxes into account does not always lead to reduced optimal turnover. The maximum fraction of capital gains which could be taxed is 100%. If optimal turnover ignoring taxes is above 100%, introducing taxes does not make a difference. This happens for some combinations of smaller fund sizes (<$1Bn) and smaller values of cost aversion coefficient ($T<30$).[17]

The tax results here are preliminary, suggesting that for a representative fund size and reasonable value of cost aversion, coefficient taxes can reconcile theoretical turnover with the data. Properly including taxes in fund strategy requires the incorporation of position-specific tax considerations at the portfolio optimization stage, in the same way as we have treated transaction costs. A full development of this line of thought is beyond the scope of this paper, however, and must await further research.

[17] Jeffrey and Arnott (1993) also make this point. They note that the marginal impact of taxes diminishes as turnover increases, disappearing at 100 percent turnover, and advise that taking taxes into account is most critical for low turnover ranges.

Trading Strategy and Turnover

Accounting for expected implementation cost leads to turnover levels which strike a balance between the timely exploitation of new information and the avoidance of excessive costs. It follows that accounting for transaction cost is an integral part of the portfolio formation process, and that trading strategy has an important place within investment strategy. In Section IV, we noted that the use of alternative price impact models leads to similar qualitative conclusions with respect to capacity and turnover, if the alternative model embodies a concave relationship between unit costs and volume. Nevertheless, trading strategy affects implementation costs. Quantitative differences in capacity, for alternative assumptions concerning trading strategy, can be substantial.

In order to illustrate this point, we now employ a more aggressive trading strategy than used in Section IV, enforcing a single-day trading horizon with respect to the implementation of an investment decision. In Figure 12 we demonstrate the effect of running a one-day value-weighted average price strategy (commonly known as one-day VWAP). We present the resulting relationship between net return and realized turnover for two combinations of fund size and cost aversion coefficients: a $5Bn fund with T=30, and a $500mln fund with T=50.

Figure 12. Net Return vs. Realized Turnover

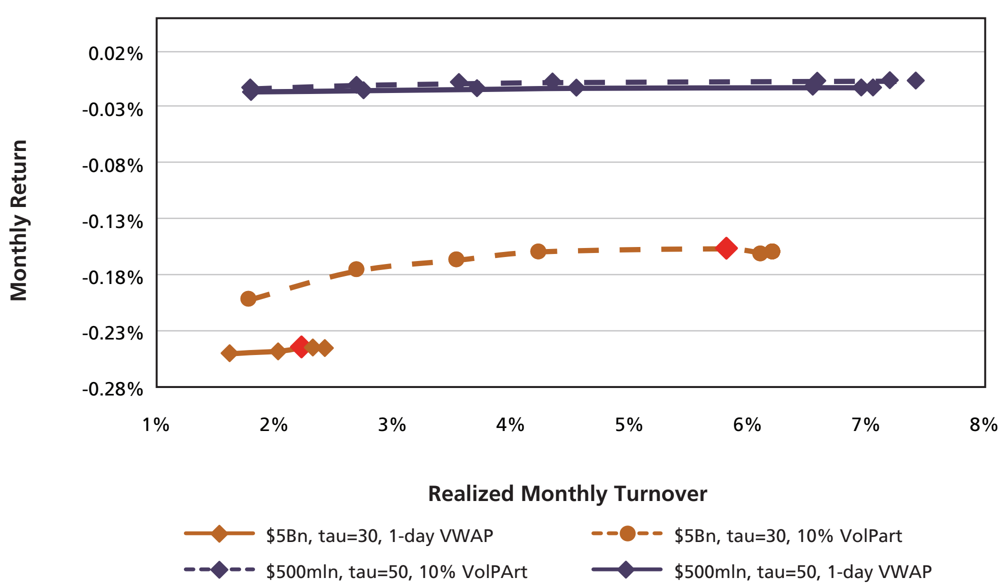

For a $5Bn fund, using the more aggressive trading strategy results in lower net returns, but also dramatically shifts the range of optimal turnover values towards zero. In fact, the magnitude of the shift exceeds the one resulting from considering taxes. We mark with red dots the optimal turnover levels for a 10% volume participation strategy (dotted line) and for a one-day VWAP strategy (solid line). The dotted line corresponds to the T=30 curve for $5Bn fund in Figure 11. The optimal turnover level for a $5BN fund and T=30 is 5.8% monthly. Taking taxes into consideration shifts this value to 4.25% per month; while running the one-day VWAP strategy shifts it even further, to 2.4% per month.

For a $500mln fund, however, the difference between running one-day VWAP and 10% volume participation strategies is negligible. Trading volumes for a smaller fund are low enough that the main component of cost is the bid-ask spread and, therefore, any trading strategy would produce a similar result.

In summary, the tactics of trading matter in the context of the overall investment strategy. The degree of importance of this observation depends itself on fund size, becoming increasingly relevant as AUM increases.

VI. Some Practical Considerations

Capacity analysis in practice typically is performed in two cases. The most common case entails a successful track record and the business desire to offer that strategy to additional clients. The second case relates to a perceived need to scale down a certain fund, and hence adjust turnover levels and execution style. In both of these cases holdings data and transaction data are usually available. Having both holdings data and transaction data benefit the analysis, but also can present difficulties when performing a capacity analysis on an existing fund.

Most of the issues stem from lack of complete data, even within a fund complex where one might otherwise expect such information to reside. At the extreme, proactive capacity analysis can be troublesome if there are no records of the investment opportunities that were considered but were not chosen to enter the holdings list. With such data, one can remove practical constraints that must be added when simulating high levels of turnover, through, for example, the use of securities that did not originally enter the trading list. Unfortunately, such data are usually not available on a historical basis, and it makes the analysis in practice limited in terms of the number of scenarios one can sensibly devise.

On the other hand, using holdings data provides information with respect to actual turnover levels, execution style, and execution patterns. Underlying transaction data allow one to calculate realized turnover as well as to identify round trip trades that cannot be observed in a time series of monthly or quarterly holdings. Round trip trades can have an effect on return calculations, and if missed for a long period of time, the cumulative affect on return can be significant. Actual transaction data also allow for the creation of daily holdings which in turn permits an accurate calculation of fund return at frequencies corresponding to trading activity

Transaction data shed light on how execution is done, including information on the typical size of an order, how long it takes to execute certain positions, and sometimes the actual execution strategy. It can also provide insight into the timing of trades, which may validate some assumptions one makes when running simulations for the fund. Examples include the type of strategy used when calculating expected cost and the most appropriate time of the month to add new transactions.

Mapping transactions from a trading cost database to portfolio holdings

Trades must be extracted from a transaction cost database and mapped to the fund holding data. The availability of a clean transaction cost data source is the first obvious issue. Funds currently using a TCA service, whether in-house or third party, pass this hurdle. This is not as big an obstacle as it may seem. Survey evidence suggests that 98 percent of large institutions and 88 percent of medium-sized institutions globally are consumers of transaction cost research.[18]

In cases where transaction cost data are not available, one must make assumptions regarding round-trip trades and turnover level within each period. One approach is to use the difference between the holdings in the beginning of the period and the holdings in the end of the period to recreate the trades that occurred within that period. This approach ignores round trip trades. It also requires additional assumptions with regard to number of round trip trades and the timing of the trades within that period, as trades in different times can yield significantly different returns. Our experience is that funds prefer to use actual fund-specific transaction cost data, and the discussion here proceeds under that presupposition.

Relevant data extracted for each transaction include ticker, decision date/time, decision shares, transaction date/time, transacted shares, side, and price. Discrepancies between the (daily) TCA data and the (monthly or quarterly) holdings information are typically reconciled by adding transaction data to resolve the discrepancy of shares held. In particular, if the end-of-month share amount for a ticker is different than the end-of-month share amount obtained from TCA transaction data for the month, transaction information may be added according to a set of heuristics. More specifically:

1) If there was an existing transaction on the correct side of the market within the month, the transacted share amount is changed to reconcile the month-end numbers.
2) If there was no transaction within the month (on the correct side), a transaction is added on the mid-month trade date for the amount of the discrepancy.

An alternative is to cull the data set down to only those data points which permit direct reconciliation. Distortion of the holdings data, however, seems to be a worse evil than ‘plugs’ for errant transactions.

Construction of a time-series of daily holdings

In order to calculate accurate returns, one must quantify the effects of all transactions on the total fund value, as well as the amount of cash held in the portfolio, and determine the total fund value on a daily basis. Once all transaction data is mapped to the fund’s portfolio holdings, a day-to-day portfolio of holdings is constructed with end-of-day share amounts and closing prices for all tickers in the portfolio, for every trading day in the analysis period.

[18] “Imperfect Knowledge: International Perspectives on Transaction Cost Analysis,” The TABB Group, March 2008.

In order to construct daily holdings, cash obtained from selling stocks is added to the portfolio's cash account. Cash used to buy stocks is subtracted from the portfolio's cash account, and all transaction costs and commissions incurred as a result of the day's transactions are subtracted from the cash account each day.

It is difficult to account for changes to the portfolio's cash account beyond those affected by the buying and selling of stock. Taking any of the month-end cash amounts specified in the fund holdings data into consideration presents problems not easily handled by the theoretical framework. In particular, the framework implicitly assumes some type of equilibrium state in which no new cash enters the portfolio, such that all trades are driven solely by the investment strategy. This can be conceptually complicated by an assumption that cash is available for purchases, which is equivalent to unlimited borrowing. We prefer this assumption over relying on actual cash flows in order to make the fund's strategy "immune" to the cash dynamics, which cannot be predicted or controlled.

Once the series of daily holdings is available, a transaction cost forecasting model is used to retrieve cost estimates for all transactions. While one of the advantages of having transaction cost data available is the ability to calculate realized costs, the demands of ex ante analysis require that one deals with simulated transactions that do not have realized costs.

Signal versus cash flow distinctions

It is useful, if not necessary, to separate trades resulting from cash flows and trades resulting from an investment decision. We call the latter *signal trades*, and the former, *cash flow trades*. The separation is done in order to analyze what we consider to be two different states of the process.

The first state, which involves the analysis of historical daily holdings, is based on signal trades only, and assumes that at each AUM level the fund is closed to new investors. This entails an evaluation solely of the investment strategy at hand. Hence, the construction of historical daily holdings based on signals is meant to evaluate the decision making of the portfolio managers and the investment strategy assuming a *static state* (no cash flow) from an AUM perspective.

The second state, which involves the analysis of historical daily holdings, including both signal and cash flow trades, assumes growth and evaluates the effects of both large injections of cash as well as signal-based trades. This second state, the *dynamic state*, entails the growth of the fund as well as the decision making process.

In our experience, the dynamic state generates higher paper returns. The difference in return is due to two factors:

- in the dynamic state, cash is reinvested in existing positions that were originally picked based on a positive signal, reinforcing that positive signal and generating higher returns;
- trades deemed as cash flow often have a relatively low percentage of median daily volume and hence a relatively low estimated cost, therefore a higher net return.

One might believe that the distinction between signal and cash flow trades is directly available from fund information sources, but this is not typically the case. In order to distinguish between a signal trade and a cash flow trade, we have developed two independent methods and vary the threshold parameters for each method. The following is a brief description of the two algorithms.

The *weight change method* takes advantage of the empirically-based assumption that cash flow trades generally introduce only a small percentage change in the weight of any given name in the portfolio relative to the total fund value. We then identify cash flow trades by comparing the end-of-day weights (taking into account the amount of the transaction decision made on that day) of any names for which a transaction decision was made to the previous day's end-of-day weight. If the weight change falls below a small threshold, the transaction decision is designated as cash flow. The threshold is a matter of opinion, not fact, but we have found ten basis points to be a useful bound. Figure 13 below illustrates an example of the frequency of signal trades and cash flow trades from a fund. It is quite clear that cash flow trades had a smaller impact on weights, involved significantly more names and had less frequency than signal trades. Cash flow trades are below the bar.

Figure 13. Signal Trades and Cash Trades

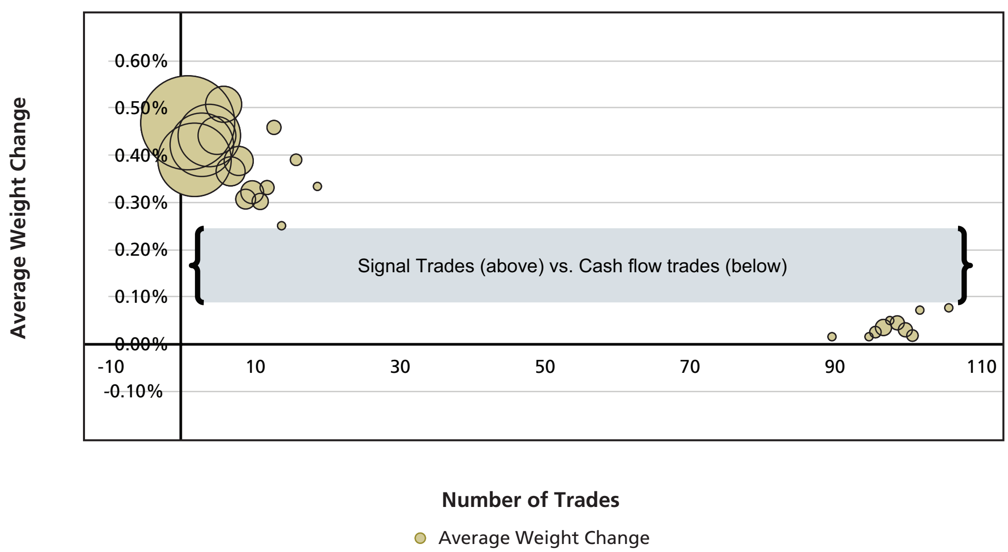

The second method exploits an alternative assumption, namely that cash flow trades are often made across a large percentage of the names in the portfolio on a given day. We identify cash flow days, as opposed to individual cash flow trades, by determining all days on which the number of names traded is greater than a threshold. All trades made on these days are deemed cash flow trades, regardless of the relative size of the trade.

Our experience with comparisons across the two methods is limited. We have found, however, that for reasonable threshold values, the percentage change in the first case and the number of names in the second, the results across methods are strikingly similar.

VI. Concluding Remarks

Capacity research has come a long way since its formal inception, most notably moving from reactive to proactive analysis. There are two new lessons to be learned from the current work.

The previous focus on turnover as a choice variable, and on its role as a proxy for implementation cost, is misplaced. Considering stock-specific transaction costs at the portfolio construction stage enables higher turnover levels, which themselves are determined through the interaction of alpha predictions and expected cost estimates. Managing the fund at a higher turnover level allows for faster processing of new information which, in conjunction with effective cost control, leads to superior net return.

The second point is closely related to the first: excluding trading strategy from portfolio construction negatively impacts returns and reduces the capacity of the fund. Lack of consideration of trading and its costs also distorts capacity analysis, and biases capacity choice downwards in general.

The use of an explicit optimization framework to reach such conclusions is both a strength and a limitation. The ability to explicate the problem formally in such a way as to isolate the contributions of individual components is part of the strength, as well as the ability to replicate and extend results. Such models also are used widely in certain segments of the asset management community, and are extensible to investment problems involving dynamic substitution of similar assets. The limitation also is fairly clear, namely that a variety of stock-picking exercises are precluded from explicit analysis.

Nevertheless, we believe that the basic lessons derived within such a framework are intuitive and more generally valid. Establishing this point is an interesting direction for future work, widening the practical value of the exercise.

References

Bogle J. (1994) Bogle on Mutual Funds: New Perspectives for the Intelligent Investor. Burr Ridge, IL: Irwin Professional Publishing.

Borkovec, M., I. Domowitz, B. Kiernan and V. Serbin (2009) Portfolio Optimization and the Cost of Trading, ITG working paper, available at http://www.itg.com/news-research-events/itg-research-trends/ITG-Paper-PortfolioOpt-031709.pdf

Bull, P., V. Serbin and H. Zhu (2009) The Capacity of Liquidity-Demanding Equity Strategies. *Journal of Portfolio Management*, Fall 2009.

Coppejans, M. and A. Madhavan (2007) The Value of Transaction Cost Forecasts: another Source of Alpha. *Journal of Investment Management*, 5(1), 65-78.

Domowitz, I., J. Glen and A. Madhavan (2001) Liquidity, Volatility and Equity Costs across Countries and over Time. *International Finance*, 4(2), 221-255.

Engle, R. and R. Ferstenberg (2007) Execution risk: It's the same as investment risk, *Journal of Portfolio Management*, 33(2), 46-59.

Grinold, R. and R. Kahn (2000) The Efficiency Gains of Long–Short Investing. *Financial Analysts Journal*, 56(6), 40–53.

ITG Inc. (2007) ITG ACE – Agency Cost Estimator, available at http://www.itg.com/news_events/papers/ACE_White_Paper_200705.pdf

Jeffrey, R. and R. Arnott (1993) Is your Alpha Big Enough to Cover its Taxes: The Active Management Dichotomy, *Journal of Portfolio Management*, 19(3), 15-25.

Kahn, R. and J. Shaffer (2005) The Surprisingly Small Impact of Asset Growth on Expected Alpha. *Journal of Portfolio Management*, 32(1), 49-60.

Karpoff, J. (1986) A Theory of Trading Volume. *The Journal of Finance*, 41(5), pp. 1069-1087.

Korajczyk R. and R. Sadka (2004) Are Momentum Profits Robust to Trading Costs? *The Journal of Finance*, 59(3), 1039¬–1082.

McDonald, I. and K. Richardson (2006) For South Korea, 'Emerging' Label Can Be a Burden. *Wall Street Journal*, July 12.

Perold, A. and R. Salomon (1991) The Right Amount of Assets Under Management. *Financial Analysts Journal*, 47(3), 31–39.

Pogue, G. (1970) An Extension of the Markowitz Portfolio Selection Model to Include Variable Transactions' Costs, Short Sales, Leverage Policies and Taxes. *Journal of Finance* 25(5), 1005- 027

SEC (2007), Registration statement of separate accounts organized as management investment companies, available at http://www.sec.gov/about/forms/formn-3.pdf

Vangelisti, M. (2006) The Capacity of an Equity Strategy. *Journal of Portfolio Management*, 32(2), 44–50.

Wermers, R. (2000) Mutual Fund Performance: An Empirical Decomposition into Stock-Picking Talent, Style, Transactions Costs, and Expenses." *Journal of Finance*, 55(4), 1655–1695.

About

ITG

Investment Technology Group, Inc., is an independent agency broker and financial technology firm that partners with asset managers globally to improve performance throughout the investment process. A leader in electronic trading since launching the POSIT® crossing network in 1987, ITG takes a consultative approach in delivering the highest quality institutional liquidity and market-leading execution services, measurement tools, and proprietary data. Asset managers rely on ITG's independence, experience, and intellectual capital to help mitigate risk, improve performance, and navigate increasingly complex markets. The firm is headquartered in New York with offices in North America, Europe, and the Asia Pacific region. For more information on ITG, please visit www.itg.com.

ITG Opt

ITG Opt®, a flexible portfolio optimization platform, helps portfolio managers develop new portfolio construction strategies and solve complex optimization problems. ITG Opt allows you to accurately model tax liability, transaction costs, and long/short objectives, while adhering to diverse portfolio-specificconstraints. With a built-in back-testing environment, open data architecture, and a user-friendly API, ITG Opt can help you more easily construct, rebalance, and back-test optimal portfolios.

ITG Data Analytics

Optimize portfolio decisions, perform transaction cost analysis, and manage risk with ITG Data Analytics™. Backed by the latest academic research and the expertise of ITG's financial engineers, ITG Data Analytics delivers real-world transaction cost estimates, high-quality data analytics, and a sophisticated set of risk models to empower decision-making and improve performance.