The Eight-Per-Cent Solution - Solving an Investment Mystery

Steve Glass - Co-CEO, Abel Noser Holdings

Past Performance Doesn't Guarantee Future Success. Or Can It?

Asset owners trying to project future investment performance know two things: No one has a "crystal ball," and past performance doesn't guarantee future success. For this reason, prudent fiduciaries look at multiple qualitative and quantitative data points to help them identify asset management firms potentially able to sustain superior returns into the future.

In theory, once evidence of a sound rigorous process (from idea-generation through trade-implementation) is identified in a manager with superior performance, investors can expect the manager to enjoy success going forward. Unfortunately, as underscored by research recently published by S&P, most managers have difficulty maintaining such success relative to their peers.

> *S&P's research quantified that for the 3-year period ending September 2019, **only 8% of US managers** who began that period in the top quartile of their respective peer group universe, remained in the top quartile at the end. Over the trailing 5-year period, the ability to maintain a performance edge was even more challenging, with less than 3% of every peer group remaining in the top quartile after five years.*[1]

Many reasons may explain why a firm, who previously outperformed their peers, might lag those peers in the future. Like good detectives, the due diligence conducted by plan fiduciaries is intended to try and identify the existence of those factors in advance.

Our research has shown that an often overlooked, or underestimated, contributor to underperformance is prohibitive trading costs. This is particularly germane in the context of manager-searches.

Not only are manager-searches instances when asset owners have to make fiduciary decisions on which manager to hire, but those managers are typically on the cusp of receiving an influx of new assets (due to their superior historical performance) that promise to make their future trades *more illiquid*.

This begs the question, to what degree will growth in a manager's assets under management ("AUM") sap their future performance (through increasingly expensive trading costs). In this regard, greater knowledge of the role trading plays in a manager's investment process can yield unique forward-looking insights.

This article is divided into three sections. *The first section provides a brief primer on what we mean by the term trading costs. The second section describes how deeper quantitative analysis of trading costs can enhance the deductive reasoning engaged in by fiduciaries trying to predict future investment success of search finalists. The final section is a case study of a manager-search involving three small cap growth managers that provides a concrete illustration of the insights possible through this due diligence.*[2]

---

[1] *S&P Dow Jones Indices Research, Does Past Performance Matter? The Persistence Scorecard, December 2019*

[2] *Acknowledgment:* Vinod Pakianathan, Managing Director played a key role in quantifying the trading cost statistics cited herein. This article could not have been completed without those efforts.

Part I. What are Trading Costs and Why are They Important?

For several decades, the trading community has recognized two fundamentally different methods of calculating trading costs: VWAP and implementation shortfall. While both methods are legitimate approaches for assessing trading acumen, they are intended to measure very different things. Recognizing the specific needs each approach is designed to address is therefore a prerequisite for correctly utilizing trading costs as a means of projecting future returns.

The goal of a VWAP trading strategy is to ensure that a manager/broker's trades are executed at prices close to the dollar-weighted average price paid for that security in the broad marketplace. The time period used in calculating the average price can be a full day, a portion of a day, or even a multi-day period.

The primary purpose of a VWAP strategy is to ensure the client's trades do not incur excessive trading costs relative to the average price paid in the broad market (for that time period). In this fashion, matching the VWAP satisfies the trader's fiduciary obligation to demonstrate the achievement of "best execution."

However, what a VWAP analysis does not do is calculate the amount of assets paid out of a fund in order to buy/sell a security. It simply quantifies the delta between a manager/broker's trade price and that security's average price.[3] The actual loss of value is calculated with the implementation shortfall approach.

In contrast to the VWAP approach, the implementation shortfall[4] and/or arrival price formula compares the execution price of a security to the price of that security at the point in time the manager first determined to buy/sell it. The arrival price cost definition is a variation of implementation shortfall, intended to measure a portion of the full implementation shortfall cost – as of the point in time a trade arrived at a particular place within the trade implementation cycle (e.g. when the trade was received by the manager's trading desk, or when the trade was received by a broker/dealer).

Put another way, implementation shortfall trading costs essentially quantify the value paid from fund assets as a result of a manager either building or unwinding a position. Given the size of institutional orders, managers sometimes "work" their orders for several days, using multiple brokers, algos and a variety of execution facilities to maximize access to liquidity and minimize market impact. And the total costs ultimately paid from fund assets to implement those orders can often exceed -100 bp.

Importantly, just as investment returns are calculated by comparing the beginning and ending values of a portfolio (after adjusting for cash flows), implementation shortfall calculates trading costs by comparing the beginning and ending prices of a trade order. This enables practitioners to better understand the portion of overall investment performance attributable to trading (in a negative way).[5] This is also one of the primary reasons why regulators mandated this calculation for complying with both MiFID II[6] and PRIIPS[7] reporting of portfolio transaction costs – in support of investor transparency, comparability and protection objectives.

---

[3] *Indeed, for larger more illiquid trades (which typically incur larger trading costs) a VWAP analysis can become a bit misleading, since the manager/broker's trade is included in the universe of trades used to determine the average price against which each trade is judged. As such, if the trade is large enough, it may become a "self-fulfilling prophesy."*

[4] *The term "Implementation Shortfall" was first coined by Harvard Professor Andre Perold in his paper, "The implementation shortfall: Paper verses reality", published in The Journal of Portfolio Management Spring 1988, Vol. 14, No. 3*

[5] *In this regard, the implementation shortfall approach is entirely consistent with the CFA Institute's definition of "Best Execution" - "The trading process most likely to maximize the value of client portfolios."*

[6] *Directive 2014/65/EU of the European Parliament and of the Council of 15 May 2014 on markets in financial instruments and amending Directive 2002/92/EC and Directive 2011/61/EU, OJ L 173/349.*

[7] *Packaged Retail and Insurance-based Investment Products (PRIIPs) - Regulation (EU) No 1286/2014*

Bottom line, while the VWAP analysis is sufficient for assessing whether a manager/broker provided best execution, it will fall short of addressing the more holistic need of calculating the impact on performance attributable to trading. To illustrate the degree to which this truism can impact a fund, Table 1 below, displays the average commission costs, VWAP trading costs, and implementation shortfall trading costs incurred by small cap growth managers over the past five years.

Table 1

Small Cap Growth Median Trading Costs[8]

	Implementation Shortfall (bp)	Commission (¢)	VWAP (bp)
2015	98	2.5	3.8
2016	125	2.4	4.0
2017	66	2.5	3.0
2018	118	2.4	4.0
2019	128	2.2	4.6

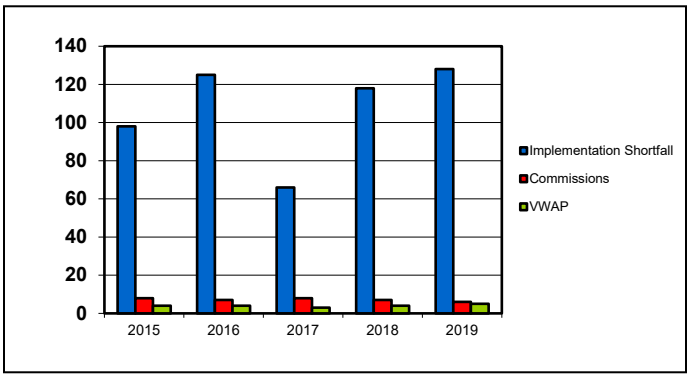

Even a cursory review shows that the VWAP and implementation shortfall methodologies yield vastly different results. This is not a surprise. It simply reflects the fundamentally different purposes of the two approaches. It also underscores the importance of using the implementation shortfall approach when determining the impact that trading has on a manager's investment performance.

In assessing that impact, it's important to recognize that the inherent nature of the manager's investment process can largely pre-ordain the level of trading costs typically incurred on their trades. For example, some managers run concentrated portfolios (with relatively few holdings) and therefore each trade, will almost by definition, be less liquid than another manager in the same peer group who holds significantly more names. Some managers employ investment strategies that call for trading into momentum (which usually entails buying/selling securities at the very moment those prices are rising/falling). Some managers employ high-turnover strategies versus a buy-and-hold approach (which will yield very different turnover rates). Some managers execute trades through certain broker-dealers in order to pay for research, and some managers have a significant number of client accounts who direct brokerage to various broker-dealers etc. (which may inadvertently result in higher trading costs).

Not surprisingly, those different trading strategies/tactics will also vary dramatically in their sensitivity to growth in AUM. Applied in a forward-looking fashion, it is the interplay between the turnover rates, and trading costs (associated with each manager's trading process), that can uncover unique insights into the degree to which growth in AUM may sap alpha.

The Manager-Search Conundrum: How to Predict Future Success?

The challenge of projecting future investment success is particularly acute in the context of manager-searches. Manager-searches represent one of the few instances where the typical asset owner (who delegates investment discretion to third party managers) makes fiduciary decisions that directly impact their fund's performance. Manager-search finalists also typically have superior past performance (they wouldn't have made it to the search finals otherwise), which exacerbates the decision-making process.

Further, the risks associated with selecting a sub-optimal finalist are compounded by the fact that the

[8] *Zeno Consulting Group, LLC: Small Cap Growth Universe*

superior historical returns which made a product attractive to one fund will likely make them attractive to other funds as well. As a consequence, the AUM of manager-search finalists may experience dramatic growth in the relatively near term.

To this end, asset owners/consultants perform extensive due diligence on search finalists. Unfortunately, this traditional due diligence has often struggled to identify criteria that can meaningfully project future performance. As a practical matter, selecting the winner therefore typically boils down to minute nuances and compatibility.

This is where a quantitative focus on trading costs may prove very useful. More to the point, we believe quantitative due diligence that projects the degree to which trading costs may increasingly sap overall returns (as AUM grows), can help minimize the risk of asset owners "chasing performance." We call this Liquidity Threshold Quantitative Capacity Analysis, or Liquidity Threshold Analysis for short.

Part II. What is Liquidity Threshold Analysis?

Ultimately, Liquidity Threshold Analysis ("LTA") is intended to demonstrate how much new money a manager can add before their current investment process reaches its alpha threshold (absent significant changes to that process, such as increasing their list of holdings, picking different types of stocks, reducing turnover rates, etc.). To be sure, it's quite possible that managers may be able to push that threshold further out by modifying their investment and/or trading process and depending on the client's investment guidelines such modifications may well be appropriate.

However, modifying an investment and/or trading process are not trivial exercises, and must be done with care and prudence. For example, adding additional names to a portfolio in order to reduce the average share size (and thereby liquidity needs) of each position necessarily requires the manager to move further down their list of preferred ideas, which means less projected alpha. That trade-off should not be taken lightly and only after thorough analysis. And even where managers successfully extend their ability to absorb additional AUM, there will always be capacity threshold, beyond which trading costs erode alpha to an unacceptable level.

Obviously, no manager accumulates assets with the intent of constraining client performance. However, at some point in time, as AUM increases, every manager reaches a trading threshold, above which they are no longer able to build or exit positions in a cost efficient manner (i.e., they have grown so large that their trade orders are too illiquid). And while not necessarily a conscious decision, the manager's interest in accumulating AUM can in effect, implicitly conflict with their client's interest in maximizing returns.

To this end, LTA simply calculates what could happen if no changes were made, and focuses on the degree to which increased AUM may result in systematic "performance headwinds" (caused by increasingly expensive trading). This exercise entails three distinct steps.

The first step is to quantify the actual loss of asset value paid from an asset owner's fund (as a result of a manager building or unwinding positions) in its current portfolio. As discussed earlier, depending on a manager's investment mandate and trading process, trading costs can range anywhere from under -20 bp for larger cap and value-type stocks, to more than -1.5% for smaller cap and growth-like stocks. For example, during the 12-month period ending Q2 2019, the trading costs incurred by Small Cap Growth strategies ranged from less than -65 bp to more than -160 bp, with a median cost of about -128 bp.[9]

[9] *Zeno AN Solutions' Q2 2019 Small Cap Growth Peer Universe*

The second step is to determine the degree to which these costs impacted overall investment performance. This is achieved by multiplying each manager's trading costs by the level of trade volume executed throughout the year (clearly a manager with low turnover can be less concerned about high trading costs, than a manager with higher turnover). These costs are imbedded in the manager's current investment performance, and from a return-attribution perspective can often represent 10-15% of a manager's gross performance (in a negative direction).[10]

For example, during the 12-month period ending Q2 2019, turnover rates among Small Cap Growth strategies ranged from less than 40% to more than 135%, with a median turnover rate of 75%. Consequently, when each Small Cap Growth manager's average trading costs were linked to their respective turnover rates, the annual impact on those managers' overall investment returns ranged from less than -100 bp to well over -2.5%.[11] With the Russell 2000 Growth index yielding an annual rate of return less than 10% over the prior five years[12], it's easy to see how trading costs can be the difference between investment success or failure.

In the context of manager-searches, this dynamic can be even more acute, since as noted earlier, manager-search finalists, almost by definition, are likely to enjoy significant near-term growth in AUM. This will result in less liquid trades, which tend to only increase trading costs. The key therefore is to understand the degree to which a manager's current trading process saps investment returns, *and the sensitivity of that process to increases in AUM.* This is the third step.

To perform this final task, the LTA engine applies a straightforward approach to calculating the ability of an investment product to absorb additional AUM without undue performance impact. It begins by using each manager's actual trading performance, as the base-line for projecting that manager's future growth in trading costs.

Once a manager's base-line trading cost profile is calculated, LTA then incrementally increases the share size of every trade executed during the review period (i.e. simulating an increase in that manager's AUM). The increased trades are then run through the LTA analytical engine to project how much the trading costs associated with those larger trades would likely increase. Those increased amounts are then added to the manager's base-line costs.

This exercise is performed on an iterative basis, reflecting increasingly larger AUM growth scenarios. At each iterative stage, the newly calculated trading costs are adjusted by the manager's historical turnover rates in order to convert those trading costs into an annualized return. That annualized trading cost return is then deducted from the manager's actual investment performance.

In this fashion LTA charts the projected erosion of each manager's original investment performance, as their AUM increases. It's important to recognize that, as explained earlier, these projections are predicated on the idea that each manager would execute the exact same trades at the exact same time as in their original base-line. LTA is therefore essentially demonstrating what might be expected had the manager done nothing different other than experience an increase in their AUM.

The following Case Study helps illustrate this dynamic.

---

[10] *This number is derived by calculating a manager's average trading costs, multiply that number by 2 (to reflect their roundtrip cost), then multiple that amount by the manager's turnover rate (to convert the aggregate cost number into an annual return). That annual trading-related return amount is then divided by that manager's gross annual investment return (i.e. the manager's net annual return plus adding back in the annualized trading cost return).*

[11] *Id.*

[12] *SPIVA, US Scorecard, Mid-year 2019.*

Part III. Case Study Overview[13]

In the fourth quarter of 2019, an asset owner conducted a search for a $150 million Small Cap Growth mandate. Three Small Cap Growth managers made it to the finals of the asset owner's search. All three finalists had experienced investment teams that had been together for long periods of time. All three finalists also had strong historical performance that outperformed their benchmark index (the Russell 2000 Growth) over the prior 3, 5, and 10-year periods.

However, the three managers had different levels of AUM, number of holdings, capitalization sizes, trade liquidity, turnover rates, and trading tactics. Further, the asset owner and its consultant were aware that a number of other large asset owners were also considering hiring each of the three finalists. Such engagements, plus the anticipated organic growth, presented a very real possibility that all three finalists could see significant growth in their AUM (of several billion dollars) over the next few years.

This raised the question of how (and to what degree) such growth might materially impact each manager's future performance, as their AUM grew. To obtain some insight into this risk, the asset owner decided to conduct a LTA on all three finalists.

Manager-Search Finalist Profiles

While all three finalists had superior 3, 5, and 10-year performance, the LTA used in the Case Study focused on the 5-year returns (i.e., a proverbial "market cycle"). As of the end of the 5-year period, the three firms exhibited the characteristics displayed in Table 2 below.

Table 2

	Manager 1	Manager 2	Manager 3
AUM	$2.84 Billion	$1.7 Billion	$2.1 Billion
Avg. # of Holdings	118	70	92
Avg. Capitalization	$3.9 Billion	$1.5 Billion	$3.5 Billion
Avg. Trade Liquidity	38.5% ADV	62.7% ADV	42.3% ADV
Avg. Turnover Rate	58%	91%	94%
Top 10 Holdings Concentration	16.8%	22.5%	21.1%
Avg. R^2	95.11	93.72	91.39
Sharpe Ratio	0.86	0.71	0.61
Annual Return	**14.1%**	**10.6%**	**12.1%**
Annual Alpha	**5.0%**	**1.5%**	**3.0%**

A review of these profiles raised a number of potential questions that the asset owner, as a prudent fiduciary, wished to investigate. These included:

- Manager 1 had the strongest historical returns, but also the most AUM. Given their robust historical returns, there was a strong possibility they would win additional searches from other asset owners. This would undoubtedly make their future trades more illiquid. On the other hand,

[13] *The manager-search described in the Case Study is conducted by a fictitious asset owner, but with real Investment manager search finalists. The historical AUM, holdings, trading activity, performance and alpha of each search finalist described herein, are the actual data (and experiences) of the three respective firms.*

Manager 1 typically held almost 70% more names than Manager 2 and 30% more names than Manager 3. This might help minimize Manager 1's trading costs relative to the other finalists, since the size of Manager 1's trades were necessarily smaller and more liquid. The asset owner therefore wanted to quantify the risk, on balance, that Manager 1's AUM might grow too big – and they would no longer be able to enter/exit their positions efficiently.

- Manager 2 had the weakest relative performance (albeit still outperforming the benchmark), but also the least amount of AUM. This suggested that Manager 2 may have more room for growth. However, they also managed a more concentrated portfolio, trafficked in smaller securities (i.e., with less liquidity), and had a fairly high turnover rate. Would the lower AUM base offset their more concentrated portfolio, illiquid securities and higher turnover rate – thereby allowing Manager 2 to grow more than Managers 1 and 3?

- Manager 3 fell in between Managers 1 and 2 in almost all categories (e.g. AUM, returns, holdings, trade liquidity, capitalization etc.). Did this represent the optimal balance, in terms of predicting the degree to which increased trading costs (due to greater AUM) might erode future performance?

Liquidity Threshold Analysis Findings

As discussed earlier, LTA begins by calculating a manager's actual trading performance, which is then used as a base-line for projecting that manager's future growth in trading costs. In this regard, during the 5-year review period the three finalists incurred average trading costs of -63 bp, -156 bp, and -153 bp respectively. These costs were a function of the actual AUM of each finalist, the characteristics of the securities they bought/sold, the market conditions in which they traded, and their respective trading tactics.

In the most recent year, Manager 1 averaged costs of -71 bp and a turnover rate of 60%. This translated into a negative impact on Manager 1's annual investment return of -85 bp. In contrast, Manager 2 averaged costs of -146 bp and a turnover rate of 79%; which equated to a negative impact on Manager 2's annual return of -2.3% bp. Manager 3 was the most expensive, with an average cost of -156 bp and a turnover rate of 94%, representing a negative impact on Manager 3's annual return of more than -2.9% bp.

In large part, these costs reflected the type of Small Cap Growth securities each finalist was trading and was a function of the size and speed (subject to liquidity constraints) with which they executed their trades. Figures 1, 2 and 3 below, detail some of the key observations of each finalist.

Figure 1

Manager 1 tended to employ a systematic trading process whereby they built positions whose prices had fallen around 50bp in the prior 2-5 days before trading, and then appreciated in value throughout the rest of the quarter.

Manager 1 was able to purchase about 60% of their orders on the first day of trading (i.e., while prices were still relatively low), and had completed about 80% by the end of day 3. The rest of the orders were traded over the following seven days.

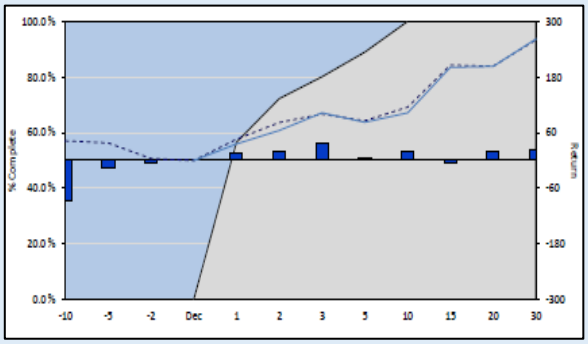

Figure 2

In contrast to Manager 1, Manager 2 typically bought securities whose prices had risen about 50 bp in the 2-5 days prior to Manager 2 beginning their trades.

While Manager 2 tended to complete their trading a little quicker than Manager 1, getting 80% done by the end of the second day and executing the residual shares over the following three days, the stock prices were already rising so Manager 2 had to pay-up in order to purchase those shares.

Further, the aggressiveness of Manager 2's trading may have resulted in price reversions following the completion of their trading (as their stock prices dipped while the index continued to rise).

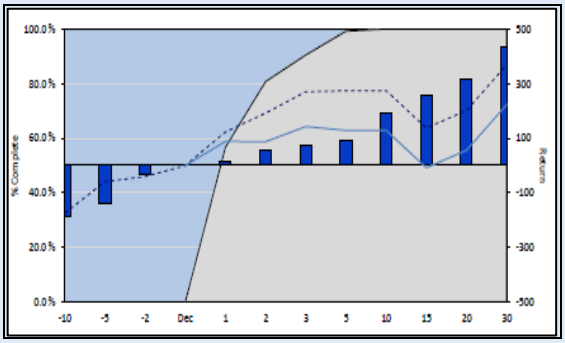

Figure 3

Manager 3 tended to trade similar to Manager 2. The securities they bought began rising about 2 days before they began trading and continued to appreciate in value throughout the entire trading period.

However, Manager 3 typically was able to buy only about 50% of their orders by the end of the first day and 65% by the end of day 2. Like Manager 1, Manager 3 often needed up to 10 days to complete their trading, which necessarily resulted in higher trading costs (as they had to buy shares at progressively higher prices).

Happily, Manager 3's slower trading appeared to avoid the price reversions experienced by Manager 2.

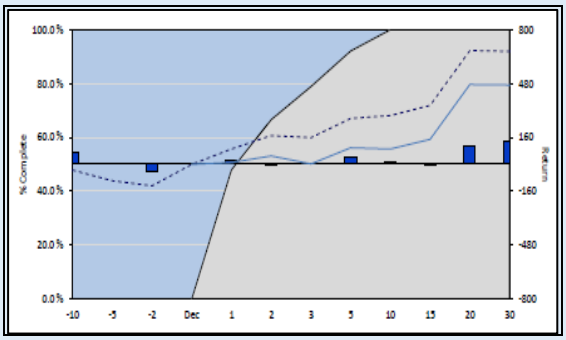

Legend

- The days before and after trading begins (the "Dec" point) are shown on X-axis.
- The gray area depicts the % traded (shown on left Y-axis) as of each respective day.
- The dotted line tracks the traded securities price movements, relative to the "Dec" date (measured on right Y-axis).

The above analysis represented the base-line upon which LTA then projected the trajectory of how each finalist's performance would decrease, as AUM (and trading costs) increased. As described in Part II, by holding everything constant except trade size, LTA estimates how much trading costs would increase if each individual trade executed during the historical period being reviewed, was incrementally larger in size.

The overall impact on annual performance is then simply a function of each manager's increased estimated trading costs along with the actual turnover rate during the historical period being reviewed. LTA then uses the resulting analysis to extrapolate what each manager's returns would have been over that historical period had the manager enjoyed progressively greater AUM (and not changed their investment philosophy).

This exercise enabled the asset owner (and manager finalists) to ascertain the point in time where growth in AUM was projected to sap performance of each finalist to an unacceptable level. Of course, each asset owner/manager needs to determine when that point has been met, based on their own respective considerations.

In the subject Case Study, the asset owner's investment guidelines called for their active managers to provide an excess return (i.e. "alpha") above their index benchmark of at least 1.5% per annum. The central question for the asset owner was therefore how much additional AUM each finalist could absorb before reaching that threshold. Figure 4 below, charts the erosion of each of the three finalists.

Figure 4

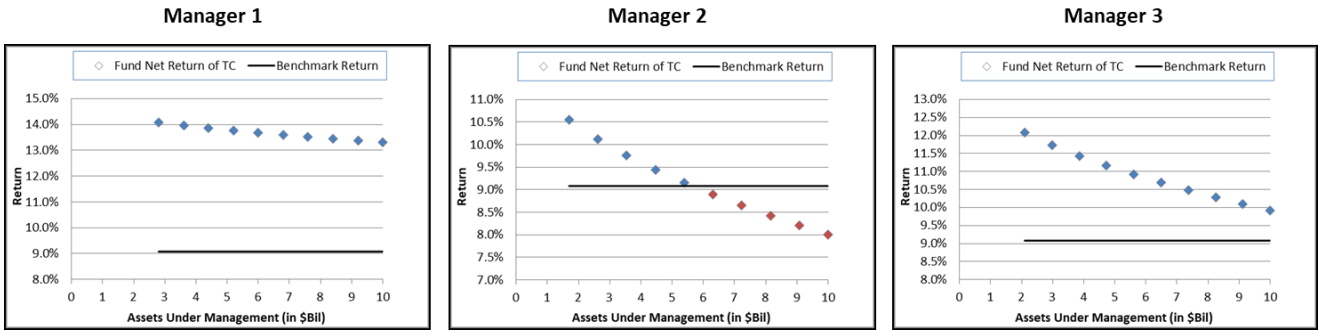

As shown in Figure 4, LTA demonstrated that notwithstanding Manager 1's greater starting AUM, the combination of slightly larger cap securities, greater number of holdings (which collectively equated to more liquid trades), and dramatically lower turnover rates, enabled Manager 1 to take on additional assets well into the foreseeable future.

In contrast, Manager 2's smaller cap names, more concentrated portfolio and high turnover rates suggested that absent a dramatic change in investment philosophy (which could itself be problematic), they were already essentially at full AUM capacity in terms of achieving the asset owner's goal of 1.5% alpha, and would lose all alpha by the time they accumulated about $5 billion in AUM.

Manager 3, while more sensitive to AUM growth than Manager 1, was nevertheless still able generate alpha even after absorbing significant additional assets. However, LTA suggested that Manager 3 would need to close its Small Cap Growth product to new investors, prior to accumulating approximately $6 billion in AUM, in order to maintain the asset owner's target alpha of 1.5%.

Conclusion - Using LTA to Help Solve the Manager-search Investment Mystery

To be clear, LTA is not a replacement for the traditional factors and considerations used to select new managers. Rather, LTA can supplement traditional qualitative and quantitative due diligence, in order to provide unique insights regarding future returns. These insights in turn enable asset owners, consultants, and managers alike, to make more informed fiduciary decisions.

All else being equal, the more efficient trading process has an inherent advantage, which the liquidity thresholds quantified by LTA either exacerbate or offset. In this context, LTA represents an additional

"touch point" for evaluating a manager's investment approach and trading aptitude and helps formulate a quantitative perspective on future returns.

Of course there may often be instances where one candidate appears to have a more efficient trading process than another, but other factors outweigh the benefits associated with that superior trading. There will be many instances where no candidate exhibits any significant concerns, but from a due diligence perspective, knowing that a particular manager's process has no "critical issues" is just as important as identifying something that warrants follow-up.

Bottom-line, all manager-searches are exercises in prudence, whose ultimate goal is to assess the likelihood of continued superior performance. We believe that when it's time to make the final decision to hire (or not hire) a particular manager, most fiduciaries will find it useful to have this information before they make that decision.

I never guess. It is an appalling habit,
destructive to the logical faculty.
Sherlock Holmes to Sigmund Freud,
N. Meyer, Seven-Per-Cent Solution, 1974

There is nothing so minute or inconsiderable,
that I would not rather know it than not.
Samuel Johnson, quoted in Boswell, Life of Samuel Johnson, 1775

LEGAL DISCLAIMER