

**The Cost of Trading Transparency:
What We Know, What We Don't Know and How We Will Know**

Gary L. Gastineau
ETF Consultants LLC
382 Springfield Avenue
Suite 206
Summit, New Jersey 07901

908-598-0440 – Office
908-598-0467 – Fax

gary@etfconsultants.com
www.etfconsultants.com

**A later version of this paper appeared in the Fall 2008 issue of
Journal of Portfolio Management under the same title**

Abstract

Results from several major studies of index composition changes indicate that trading transparency accounts for more than half the market impact cost of composition changes in the two most popular U.S. indexes. The growth of transparent custom index ETFs and the introduction of actively managed ETFs with relatively transparent trading will increase investor exposure to trading transparency costs. Fortunately, researchers can quantify the performance penalty that trading transparency imposes on investors.

The introduction of the first actively managed exchange-traded funds (ETFs) in the United States is an appropriate time to examine the cost of transparent trading to fund investors. These new, limited function actively managed ETFs must reveal portfolio changes before the market opening on the day after the trades are made. This is not a massive change, but it is significant.¹ Hougan [2008] correctly concludes that reducing these funds' trading transparency or, in his words, reducing their "front-running" costs is the primary motivation for this change. We expect this fund model to see extensive use and we will have more to say about these limited function actively managed ETFs below. However, any discussion of the cost of trading transparency must begin with index funds, whose shareholders are the greatest victims of trading transparency costs.

The S&P 500 and the Russell 2000 – the two most popular templates for indexed portfolios in their respective capitalization ranges – are well-known U. S. indexes with thoroughly studied trading transparency costs. The use of "fundamental" indexes, back-tested custom indexes and benchmark indexes from publishers other than S&P and Russell as templates for exchange-traded funds (ETFs) has created additional opportunities for traders to profit from transparent changes in indexed portfolios. A characteristic of all indexed portfolios is that index changes are published *before the portfolio manager trades to implement the change*. At present, most of the index changes for newer ETFs have a trivial market impact relative to the S&P and Russell index changes, but changes in a few of the more successful new ETF's indexes are starting to appear on index composition arbitrage lists.

In nearly all actively managed mutual funds and other actively managed portfolios, trading plans are considered highly confidential.² Absent confidentiality, advance disclosure of portfolio changes could be misused by traders who value any information that might help them scalp a profit by anticipating or participating in someone else's trading activity.³

Revelation of any portfolio trading plans in advance of execution discloses a specific demand for liquidity from securities markets. Publishing this demand for liquidity reveals information that thoughtful fund portfolio managers and investors have generally preferred *not* to have disclosed. If liquidity demands are revealed only by the entry of anonymous orders that can be executed over a number of days or weeks, the market impact of these liquidity demands is often modest. However, if liquidity demands are openly published and specified as to source, size and time, the market impact cost of a transparent trade is likely to be significantly greater than the cost of a comparable anonymous trade.

With increasing use of transparent trading by hundreds of exchange-traded index funds, it is inevitable that multiple transparent traders will sometimes announce plans to buy or sell the same security during the same narrow time window. Once similar trading intentions have been announced at the same time by several of these funds, it will be difficult for any of the funds to change course or delay the trade. If these trades were not disclosed, one or more of the managers could decide to trade patiently – or not at all if stock prices changed too much to make the trade attractive. One problem with changing an announced trading plan to reduce the cost of transparency is that market participants will rely on published trading intentions. Indeed, a transparent trading product structure seems to create public entitlement to knowledge of trading plans.

To measure the cost of trading transparency to investors, we need to separate the cost of *transparency* from the cost of *liquidity* in transparent trades. To this end, it is useful to define the cost of trading transparency as *the total trading cost of implementing a fully or partially transparent transaction minus the liquidity cost of an otherwise comparable non-transparent transaction*.⁴ The price changes associated with transparency in benchmark index portfolios usually precede the heavy trading volume associated with liquidity demand; but some transparent trades do not feature a time lag between determination of the change and the beginning of its implementation. Our definition focusing on cost differences covers all cases. Later sections and endnotes will clarify the measurement of specific differences.

How Trading Plans Become Transparent

The transparent trading process of ETFs using custom indexes and other pre-announced trades should give us enough information to isolate the incremental cost of a number of kinds of trading transparency to fund investors. Before we attempt to measure transparency cost, however, we need to examine some of the ways that trading plans become transparent. Exhibit 1 lists major categories of transparent and semi-transparent trades. The *type* of transparency characteristic of an index or an active management investment process determines how easily we can isolate and

measure the cost of its trading transparency. Some index and trade disclosure policies make the cost of transparency easier to isolate than other disclosure policies.

(Insert Exhibit 1 about here)

Benchmark Index Composition Changes

In the early days of indexing, the expectation of indexing advocates (voiced by Malkiel [1973] pp. 226–7 and Samuelson [1974] among others) was that passive “indexing”⁵ would reduce transaction costs because portfolio turnover would be as low as possible. In the context of this expectation, the magnitude of the transaction costs regularly experienced by investors in portfolios based on some major indexes is ironic in the extreme.⁶

In fairness to the publishers of benchmark indexes, they have been sensitive to the impact of transaction costs on index fund investors. Of course, index publishers cannot eliminate all composition changes. Companies are acquired or go out of business and new companies sell shares to the public.⁷ Only the S&P 500 and the Russell 2000 have had enough committed indexed assets for their composition changes to generate much excitement. The cost of turnover in even these indexes has not been an issue outside the “pretty sleepy community”⁸ of index publishers and their clients until recently. Now, the fact that transparent index composition trades are often costly is well understood.⁹

There are two very distinct types of benchmark index composition change transparency:

(1) Most major benchmark index families are based on rules. The rules for indexes in the Russell, MSCI, FTSE and Dow Jones index families are well known to index composition arbitrage traders.

(2) In contrast, Standard & Poor’s has an Index Committee that makes decisions to change the composition of its best known indexes in secret and announces the changes, after the market close, a few days before the change becomes effective.

Looking first at Standard & Poor’s, the S&P Index Committee has a number of formal policies;¹⁰ but the Committee’s selection of a replacement company is rarely so inevitable that traders will mark up the share price of a replacement candidate on the chance that it will be selected by the Committee on a specific date. David Blitzer, Chairman of the S&P Index Committee, has estimated that there are probably 50 or more companies not in the S&P 500 that meet the criteria for inclusion at any moment.¹¹ It is usually safe to assume that the market impact of an S&P

index addition begins when it is announced. Deletions from S&P indexes are usually easier to predict than additions – and typically have much less market impact.¹²

Dow Jones, FTSE, MSCI, Russell and other rules-based index publishers sometimes have their own index committees, but their committees have far less scope and influence than the Standard & Poor's Index Committee. Anyone familiar with these index publishers' rules can anticipate most index changes in advance of any formal announcement or the arrival of a reconstitution date. An important characteristic of most of these indexes is that the transparency of future composition changes builds as the probability of a specific change approaches certainty. Because the probability of a specific change gradually approaches 100% over a number of months, it is more difficult to choose a time to start measuring transparency costs than in the case of transparency by press release that characterizes S&P index changes.

As the market impact of rules-based index composition changes has become increasingly apparent – and disruptive – index publishers have modified their rules to reduce the number of index composition changes, spread the changes over time and make the market impact of changes harder for analysts and traders to measure. In an attempt to tame its reconstitution beast, Russell recently inserted a buffer zone between the Russell 2000 and the Russell 1000 to reduce movement between those two indexes.¹³

Chen, Noronha and Singal [2006] (CNS [2006]) measured the *liquidity* cost of Russell additions or deletions from the time the additions or deletions were fully determined by the Russell rules for 1990 through 2002. The CNS [2006] calculations of the cost of index composition changes capture the cost of liquidity, but they *do not include any significant element of the cost of transparency in the Russell changes because the changes were fully determined by the date CNS [2006] began to measure the composition change effect.*

Relative to announced S&P index changes, the trading transparency cost of Russell index changes is difficult to measure accurately or directly. The Russell rules-based index changes become certain gradually as the end of May approaches. While the changes are not certain until the end of May, a moment's reflection will suggest that the number of stocks whose status is significantly in question will be very small sometime before the composition changes are official.

The effect of transparency cost is relatively easy to measure in S&P index additions. A change is 100% certain when it is announced by S&P. We need not be concerned about anticipatory trades because few traders speculate on the S&P Committee's choice before the announcement. Using

S&P 500 addition data from October 1989 through December 2002, CNS [2006] found that the abnormal return for stocks added to the S&P 500 on the first trading day after the announcement of their selection (+5.12%) was over 60% of the total abnormal return from the closing price just before the announcement until the close on the effective day of the index change (+8.37%). The trading session on the effective day is the time of greatest trading volume (liquidity demand), but the price change on the first day of trading after the announcement is usually the largest daily price change. Trading volume on the first day after the announcement is usually much larger than volume on an average trading day, but that day's volume is typically much smaller than volume on the effective date when the demand for liquidity is manifest. If we take the first day after the announcement's abnormal return as a reasonable proxy for an S&P 500 addition's transparency cost, observed price behavior suggests that transparency cost has been more than 60% of the total market impact of S&P 500 additions through the moment the change becomes effective.

While CNS [2006] did not attempt to measure the transparency cost reflected in gradual price changes over the period *before* the Russell 2000 additions and deletions were fully determined, Madhavan [2003] measured at least *part* of the Russell transparency cost. Using Russell reconstitution data for 1996 through 2002, Madhavan calculated differential returns of Russell 2000 addition and deletion portfolios for the three months (March through May) *before* the changes were fully determined¹⁴ and the returns for June – *after* the index changes were fully determined and the demand for liquidity caused price changes through the end of June.¹⁵ Madhavan found that the June return (+10.78%) was less than 40% of the total March through June cumulative return of the differential portfolios (+29.04%).

These two very different indexes and different analytical methodologies suggest that more than 60% of the measurable cost of S&P 500 and Russell 2000 composition changes has been attributable to either the sudden arrival (S&P) or the growing certainty (Russell) of information on index changes – to the transparency of pending trades.¹⁶ This evidence indicates that the transparency cost element is greater than the cost of the liquidity needed to support the actual trading undertaken to implement changes in portfolios tracking these indexes.

These index studies certainly indicate that the cost of transparency is greater than many observers have believed, but the very large, high profile S&P 500 and Russell 2000 index changes may not be fully representative of smaller transparent trades. It is also possible that the price changes we attribute to transparency in these index changes tend to reduce or offset price changes that would otherwise *increase* the costs we attribute to liquidity demand. To state the possible effect in another way: Price changes occurring in advance of full determination of an impending change may reduce the apparent liquidity cost. We need data on more diverse trades

to reach a conclusion on the cost of transparency in a wider range of transparent trading situations.

S&P's announcements of index changes based on a non-transparent committee selection process are easier to evaluate for transparency cost than the rules-based procedures of most other index publishers that reveal the growing certainty of changes as the determination date approaches. The trading associated with the Russell reconstitution is largely invisible to the mainstream financial community, but the probabilities of specific Russell composition changes are carefully evaluated by participating traders. If index composition arbitrage players can profit by monitoring other rules-based indexes, upcoming composition changes in those indexes will also be anticipated by changes in the market prices of affected stocks in advance of full certainty of the index change. However, given the complexity of measuring the degree of trader anticipation of rules-based changes, the S&P-style "press release" or announcement disclosure model is likely to be our most useful source of transparency cost data.

ETF Custom Indexes

Issuers of new exchange-traded funds after 2001 had little choice but to base their products on less popular benchmarks or on custom indexes designed specifically for the ETF market. In most cases, these indexes have no significant applications other than as a fund template for a single ETF. This singularity of application suggests the possibility of relatively uncomplicated transparency cost analysis.

Most of the new custom index ETFs are small, but some new funds that have captured investor imaginations have attracted more than a billion dollars in assets. The use of index weighting policies that are not based on shares available to trade will increase trading costs for some small-cap issues entering or leaving these funds. The more the portfolios of these ETFs depart from the standard of capitalization weighting with a float adjustment, the greater the demand for liquidity will be whenever a composition change includes transactions in less liquid securities.¹⁷ A mix of relatively larger transactions in small-cap stocks and smaller transactions in large-cap stocks will provide a diverse database for trading transparency cost measurement.

If a custom *index process* or a transparently traded ETF *investment process* is opaque – with trading transparency achieved by announcement – the cost of the trading transparency can be readily measured relative to the cost of a comparable non-transparent trade. Pre-trade transaction cost models provide good estimates of average transaction costs if a trade is executed efficiently. These pre-trade cost estimates are developed from data on millions of anonymous trades. A

“press release” transparent fund that consistently experiences higher transaction costs than pre-trade transaction cost models predict will give us a useful measure of the incremental cost of trading transparency¹⁸ to the fund’s investors.

A database of S&P-style press release transparent trades is already being accumulated with each composition change in custom index ETFs that do not disclose their index rules. The cost differences between transparent trades initiated by changes in these indexes and comparable trades made anonymously in other portfolios will enable us to measure the added cost of transparency for the ETF’s trades.¹⁹

Some of the custom indexes used for ETFs are similar to the Russell indexes in the sense that the rules are published and anyone who cares can determine the changes before they are officially announced, simply by applying the rules. The extent to which index arbitrageurs attempt to front-run these funds’ transparent portfolio changes is more difficult to measure than measuring the transparency cost of press release disclosures of portfolio changes. While an analysis along the lines of Madhavan [2003] could be used to measure the cost of this rules-based transparency, there is no reason to believe that observations from press release portfolio changes will not provide a reasonable estimate of transparency costs for less readily measurable rules-based portfolios as well. Clearly, a custom index ETF that has a few billion dollars in assets and a portfolio of stocks with small average capitalizations will soon become known to and loved by arbitrageurs, however its trading plans are disclosed.

For readers interested in learning more about fund transaction costs and transaction cost measurement, Edelen, Evans and Kadlec [2007] is useful and thought-provoking. Extensive background publications listed on the ITG website, www.itg.com, explain both the role of pre-trade cost analysis models and the calculation of actual trading costs.

Actively Managed ETF Trades

Analysis of actively managed ETF trading transparency will be similar to the analysis of a press-release announced custom index ETF trade. The initial actively managed ETFs will make most of their portfolio changes in a single trading session until the fund becomes too large to make such trades efficiently. If a particularly large trade is to be made in a fund with an active management process, the transaction might be announced piecemeal and executed in tranches over several days. However disclosure of an actively managed ETF trade or series of trades is handled, the cost of transparency can be calculated by subtracting the pre-trade cost estimate of the most efficient way of executing the entire trade anonymously from the actual cost of the transparent or semi-transparent trades. This is the cost of the greater transparency (i.e., greater

than that of an actively managed mutual fund) that results from the requirement that each day's portfolio changes be revealed before the next day's trading begins.

A secondary issue affecting measurement of the cost of transparency in these actively managed ETFs is that an active manager's trade is not "informationless".²⁰ In an actively managed fund with a portfolio selected by a quantitative model that has stock-picking capabilities or by a human portfolio manager who has skill, a trade has greater future fund performance implications than an index composition change. The evidence that there is value in the purchase and sale decisions of at least some mutual fund portfolio managers (Cremers and Petajisto [2007] and Green and Hodges [2002]) is very persuasive, as is the evidence that some investors can identify these managers in advance of such superior performance (Harlow and Brown [2006], Wermers, Yao and Zhao [2006] and Kacperczyk, Sialm and Zheng [2006]). In this context, front-running of a good manager's trades may hurt the performance of the good manager by *more* than the cost of informationless transparency in an index fund trade. On the other hand, the absence of arbitrage activity around an actively managed fund's transparent trades will suggest that the market has low expectations for the manager or management process.²¹

Effect of Transparency Costs on Fund Performance

Generalizations in advance of collecting data are often called hypotheses. With appropriate trepidation I will hypothesize the range of trading transparency costs various kinds of ETFs might incur.

The cost of trading transparency is unlikely to be a major concern for every custom index or actively managed ETF. Some funds will make very few portfolio composition changes and some small actively managed funds that trade before disclosing a change will be able to make most trades efficiently in a single day before the change is revealed.

Because actively managed ETFs will hold and trade securities with highly variable liquidity characteristics and will change the composition of their portfolios at highly variable rates, the expected cost of a fund's portfolio composition trading depends on its specific costs of transparency, on the liquidity of the securities the fund trades and on the annual percentage change in the composition of the fund's portfolio. Trading transparency is costly because traders attempting to exploit transparency anticipate a profit when subsequent liquidity demand causes additional market impact. The logic behind trading in front of a transparent trade is that any front-running profit will be roughly proportional to the liquidity cost of the anticipated trade. Thus, the expected cost of anonymous liquidity trading is a logical base for transparency cost analysis to build on.²²

(Insert Exhibit 2 about here)

Exhibit 2 illustrates a range of anonymous trading liquidity costs associated with two hypothetical funds. The first fund is a large-cap fund. With minimal (5% per year) portfolio changes, the fund's expected liquidity transaction costs (two-sided) are just 3.3 basis points annually. With 50% composition changes, the annual liquidity trading cost for the large-cap fund rises to 33 basis points. The second fund is a small-cap fund. With 5% annual composition changes each year, two-sided liquidity costs for that fund are 13.1 basis points, rising to 131 basis points with 50% two-sided composition changes. The liquidity trading cost assumptions are the sum of spread plus price impact from Table III of Edelen, Evans and Kadlec [2007] for the large-cap and small-cap funds in their sample. Commissions are not included. These costs are very close to Plexus Group data for similar institutional trades.²³

Only additional research on the transparency cost of ETF portfolio trades can pin down the relationship between transparency and liquidity costs. If the cost of transparency relative to the cost of liquidity is as great in these ETFs as the S&P 500 and Russell 2000 data suggest, the transparency cost will be more than 1.5 times the liquidity cost *of the trades that are transparent*. If the 1.5 times multiple is approximately correct, the cost of trading transparency in many indexed and actively managed ETFs will be greater than many of these funds' cash expense ratios and it will be highly variable from fund to fund.²⁴

Fund Portfolio Scale Trades

Portfolio scale trades in actively managed mutual funds, the last of the trade types listed in Exhibit 1, are "semi-transparent." As the size of a non-indexed mutual fund increases or decreases, the manager rarely changes the size of all fund positions proportionately, but cash inflows and outflows require regular trading in many of the fund's holdings. The most useful studies measuring the cost of mutual fund flow trading (Edelen [1999] and Edelen, Evans & Kadlec [2007]) conclude that the portfolio trading costs of accommodating investor flow into and out of a mutual fund are the largest component of trading costs for many mutual funds. The cost of flow trading in the mutual funds Edelen, Evans & Kadlec [2007] studied averaged 75 basis points (0.75%) annually.

So far, ETFs in the U.S. have largely avoided imposing flow trading costs on their ongoing shareholders. ETF in-kind creation and redemption costs are borne initially by ETF market makers and, ultimately, by the market makers' secondary market trading customers, not by ongoing shareholders of the funds. Any market impact cost to ETF shareholders of trader flow

into and out of ETFs has been very modest.²⁵ However, the principle of protecting ETF shareholders from the cost of flow is not a feature of every existing ETF or every proposed ETF.

Conclusion

The information we have today on the cost of transparent trading comes almost entirely from index composition changes and two very limited experiments with trading transparency. The largest of these experiments was “sunshine trading,” which was used briefly in connection with dynamic, liquidity demanding portfolio insurance transactions prior to the 1987 market break.²⁶ Of these data sets, index composition changes have been studied most intensely and over several decades; but this transparency “experience” does not give us the kind of data we need to measure the cost of trading transparency in smaller trades.

The growth of custom index ETFs and fully or partially transparent actively managed ETFs should generate useful data on trading transparency costs. Until the database of such trades is larger and covers a longer period, the only reasonable working hypothesis is that trading is costly and that transparent trades are significantly more costly than nontransparent trades.

What is the significance of trading transparency costs to investors? Most importantly, investors must realize that indexing is a low cost way to obtain market matching returns *only if the template index for a fund does not have significant embedded trading transparency costs*. In addition, investors and their advisors must evaluate funds in a new way. The focus of most fund analysis on expense ratios is misplaced. Trading costs for many index funds, especially when we include the costs of trading transparency, are greater than the funds’ disclosed expenses.

Exhibit 1 – Transparent Trades and Their Market Impact

Type of Trade Characteristics	Benchmark Index Composition Changes	ETF Custom Index Fund Composition Changes	Actively Managed ETF Trades	Fund Portfolio Scale Trades
Speed of Implementation	Revealed gradually (rules-based) or announced in advance. The official change is at a precise moment.	Revealed gradually (rules-based) or announced. Implementation in a single trading session or less.	Single trading session	Continuous
Ease of Market Impact Measurement	Slightly more difficult as index publishers change methodologies.	Varies with index methodology. May be easy to measure.	Depends on fund policy. May be easy to measure.	Systematic measure would be of limited value.
Market Impact	Function of the popularity of the index as a fund template. Many studies.	No systematic studies yet. Largely a function of trade size and stock float.	No systematic studies yet. Largely a function of trade size and stock float.	Few definitive studies. See Edelen, Evans and Kadlec [2007].

Exhibit 2 Annual Liquidity Transaction Cost Estimates
(basis points)

	(1) Spread and Price Impact (One-Side)	(2) 5% Portfolio Change .05 x Two Sides x Col (1)	(3) 50% Portfolio Change .50 x Two Sides x Col (1)
Large-Cap Fund	33	3.3	33
Small-Cap Fund	131	13.1	131

Source: Edelen, Evans and Kadlec [2007]. See text.

BIBLIOGRAPHY:

Almgren, Robert and Neal Chriss, “Optimal Execution of Portfolio Transactions,” *Journal of Risk*, 3, Winter 2000/2001, pp. 5 – 39.

Arnott, Robert, D., Jason Hsu and Phillip Moore, “Fundamental Indexation,” *Financial Analysts Journal*, March/April 2005, pp. 83 – 99.

Bell, Heather, “Straight From the Source: John Prestbo,” *Index Universe.com*, August 30, 2007, http://indexuniverse.com/index.php?option=com_content&view=article&id=3081&Itemid=29

Bernstein, Peter L., *Capital Ideas, The Improbable Origins of Modern Wall Street*, Free Press, 1992.

Blume, Marshall E. and Roger M. Edelen, “S&P 500 Indexers, Tracking Error and Liquidity,” *Journal of Portfolio Management*, Spring 2004, pp. 37 – 46.

Broms, Todd J. and Gary L. Gastineau, “The Development of Improved Exchange-Traded Funds in the United States,” in Fuchita, Yasuyuki and Robert E. Litan, *New Financial Instruments and Institutions*, The Brookings Institution, 2007, pp. 193 – 209, and in *A Guide to Exchange-Traded Funds and Indexing Innovations* (Sixth Edition), Fall 2007, Edited by Brian R. Bruce, pp. 16 – 26

Chacko, George C., Jakub W. Jurek and Erik Stafford, “The Price of Immediacy,” *Journal of Finance*, June 2008, pp. 1253-1290.

Chen, Honghui, Gregory Noronha and Vijay Singal, “Index Changes and Losses to Index Fund Investors,” *Financial Analysts Journal*, July-August 2006, pp. 31 – 47.

Cremers, Martin and Antti Petajisto, “How Active Is Your Fund Manager? A New Measure that Predicts Performance,” Working Paper, January 15, 2007.

Edelen, Roger M., “Investor Flows and the Assessed Performance of Open-End Mutual Funds,” *Journal of Financial Economics* 53, 1999, pp. 439 – 466.

Edelen, Roger M., Richard Evans and Gregory B. Kadlec, “Scale Effects in Mutual Fund Performance: The Role of Trading Costs,” Working Paper, March 17, 2007.

Gastineau, Gary L., “Comments of Nuveen Investments on Actively Managed Exchange-Traded Funds, File No. S7-20-01,” January 14, 2002.

<http://www.sec.gov/rules/concept/s72001/gastineau1.htm> .

Gastineau, Gary L., “Is Indexing Ready for the 21st Century, Parts I and II,” *Journal of Indexes*, April/May 2006a and June/July 2006b.

Green, Jason T. and Charles W. Hodges, “The Dilution Impact of Daily Fund Flows on Open-End Mutual Funds,” *Journal of Financial Economics*, vol. 65, July 2002, pp 131 – 158.

Harlow, W.V. and Keith C. Brown, “The Right Answer to the Wrong Question: Identifying Superior Active Portfolio Management,” *Journal of Investment Management*, Fourth Quarter, 2006.

Hougan, Matthew, “It’s All About Front-Running,” IndexUniverse.com, April 14, 2008,

<http://www.indexuniverse.com/component/content/article/31/3993.html?Itemid=19> .

Kacperczyk, Marcin, Clemens Sialm and Lu Zheng, “Unobserved Actions of Mutual Funds,” Forthcoming in *Review of Financial Studies*, 2006.

Khandani, Amir E. and Andrew W. Lo, “What Happened to the Quants in August 2007?”

Journal of Investment Management, Fourth Quarter 2007,

http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1015987

Madhavan, Ananth, “The Russell Reconstitution Effect,” *Financial Analysts Journal*,

July/August 2003, pp. 51 – 64.

Malkiel, Burton, *A Random Walk Down Wall Street*, 1st Edition, 1973, 9th Edition, 2007, W.W. Norton.

Murguía, Alejandro and Dean T. Umemoto, “An Active Approach to Indexing,” *Journal of Financial Service Professionals*, January 2006, pp. 1 – 8.

Prestbo, John A., Ed., *The Market’s Measure*, Dow Jones, 1999.

Rudolph, Sarah, “The Russells Move to ICE,” *Futures Industry*, September/October 2007, pp. 46 – 48.

Samuelson, Paul A., “A Challenge to Judgment,” *Journal of Portfolio Management*, Volume I (Fall), 1974, pp. 17 – 19.

Scism, Leslie and Ian Salisbury, “Missing the Mark,” *The Wall Street Journal*, March 3, 2008. p. R4.

Standard & Poor’s, *Innovation and Evolution – The S&P 500*, [2007].

Treynor, Jack, “Why Market-Valuation-Indifferent Indexing Works,” *Financial Analysts Journal*, September/October 2005, pp. 65 – 69.

Wermers, Russ, Tong Yao and Jane Zhao, “The Investment Value of Mutual Fund Portfolio Disclosure,” Working Paper, October 2006.

Endnotes:

The author thanks Don Chance and Vijay Singal for valuable suggestions and comments.

¹ This approach to actively managed ETFs was described in Gastineau [2002]. Like the indexed ETFs that are the principal subject of the present paper, this actively managed ETF structure is vulnerable to transparency costs that increase as the average liquidity of portfolio securities decreases and as fund assets increase.

² The SEC mandates a Code of Ethics for employees of Registered Investment Advisors under Rule 17j-1 of the Investment Company Act of 1940. Among the requirements of this rule is detailed reporting of securities transactions by Access Persons and Advisory Persons who have knowledge of fund transactions.

³ Fully transparent actively managed funds are not without precedent. In late 1999, Don Luskin, a highly regarded former vice chairman of Barclay's Global Investors, co-founded MetaMarkets. MetaMarkets launched OpenFund, a mutual fund so transparent that its portfolio management discussions and trading activities were broadcast over a live webcam. The market break which began in 2000 was probably more responsible for the OpenFund's lack of success than its transparency, but the fund did not attract many assets even before the market break.

⁴ If transparency "rules" suggest a point-in-time execution (e.g., at a market close for index calculation purposes), the comparable non-transparent trade typically would be a more patient trade that might use a VWAP or other algorithm. Only if the trade was relatively small might the comparable non-transparent trade be a market-on-close transaction for the entire position.

⁵ The first indexed portfolio (created in 1971) was not based on an index. It was designed "to hold an equal dollar amount of each of the 1,500 or so stocks listed on the New York Stock Exchange, which seemed the most appropriate replication of 'the market'." (Bernstein [1992] p. 247.)

⁶ No one could have anticipated in the mid-1970s that indexing would be so successful that index changes would lead to index fund trading costs far greater than the trading costs in an actively managed fund with portfolio changes similar in size but less transparent to other traders.

⁷ To illustrate the inevitability of index changes, General Electric is the only company now in the Dow Jones Industrial Average that was in the DJIA when it was first calculated in 1896 (Prestbo [1999], p. 11).

⁸ Bell [2007] cites this characterization by John Prestbo of Dow Jones Indexes.

⁹ Two excellent recent studies of index composition change costs are Blume and Edelen [2004] and Chen, Noronha and Singal [2006]. Murguía and Umemoto [2006] and Gastineau [2006a and b] describe these and other studies.

¹⁰ The Index Committee occasionally ignores one or more of these policies when it makes an index change.

¹¹ Standard & Poor's [2007], p. 6.

¹² Part of the reason S&P 500 deletions have less market impact is that most of them are the inevitable result of a change in the company's status that is obvious before the Index Committee makes the change official. Only uncommon deletions like S&P's decision to drop foreign companies from the S&P 500 in 2002 or an occasional house cleaning to eliminate stocks that no longer "reflect and represent the U.S. stock market" are difficult to anticipate.

¹³ Most other rules-based index publishers use buffer zones as well.

¹⁴ To clarify, Madhavan constructed and calculated returns for portfolios that were long the ultimately determined Russell 2000 additions and short the ultimately determined Russell 2000 deletions beginning three months before market participants could be certain what the ultimate additions and deletions would be. The prices of these stocks generally moved up (additions) or down (deletions) as their status became more clear. In contrast to S&P deletions

which usually coincide with a corporate event, stocks are removed from the Russell 2000 in the annual reconstitution solely on the basis of an increase or decrease in their capitalization ranking.

¹⁵ The official reconstitution – and the heaviest demand for liquidity – occurred at the end of June during these years.

¹⁶ The differences in methodology are not disturbing because the authors of these papers are focused on market impact costs associated with index composition changes. Madhavan [2003]’s use of a three-month window to measure what we characterize as the market impact cost of trading transparency is necessarily arbitrary, but reasonable.

¹⁷ On average, the composition change transaction costs for these index funds will be greater than for a cap-weighted fund, slightly offsetting some of the expected performance advantages described in Arnott, Hsu and Moore [2005] and Treynor [2005]. Examples of this phenomenon were noted in Scism and Salisbury [2008].

¹⁸ Pre-trade transaction cost estimation models assume that a trader uses an execution strategy designed to minimize transaction costs. If a transparent trade must be executed within a narrow time window or if procedural requirements for an ETF portfolio change otherwise require a sub-optimal execution strategy, the portfolio will often experience higher trading costs than it would incur with an efficiently executed anonymous trade. Such requirements undoubtedly account for some of the index change costs in the S&P 500 and Russell 2000 experience. Although the additional cost may not be solely a result of revealing a demand for liquidity, the additional cost is part of the transparent trade execution process.

Almgren and Chriss [2000] supports relatively aggressive trading tactics even when trading is anonymous, but many portfolio managers practice patient trading because they believe they improve returns by providing rather than demanding liquidity. Khandani and Lo [2007] describe a historically profitable hedge fund trading strategy that provides liquidity to aggressive traders, suggesting that hedge funds have often profited from the fact that patient trading generally reduces transaction costs and improves returns.

¹⁹ The random arrival of orders makes actual transaction costs highly variable in specific instances. This variability requires analysis of a large number of transactions and calculation of average costs.

²⁰ The addition or deletion of a stock in a benchmark index or a custom index designed as the template for an exchange-traded fund carries no information other than that the index composition is being changed and, as a consequence of the index change, a stock will be added to or deleted from one or more portfolios. The fact that some of these indexes are called active, dynamic, fundamental or intelligent indexes has not been a reason to believe that the purchase by an index ETF of a security added to its index has significant long-term implications for the performance of that security.

²¹ Frontrunning of a fund’s trades may capture much of the value the manager of a transparently trading fund brings to her stock selections. The presumed glamour of being an active ETF pioneer would be at least partly offset by the realization that trading transparency is a flawed approach to any form of active management. While my personal expectation is that a skilled portfolio manager will be reluctant to accept the assignment of managing a fully transparent actively managed ETF, some managers may reach a different conclusion. In any event, we will be able to measure the cost of these funds’ trading transparency to the funds’ investors.

²² While strict proportionality between transparency cost and liquidity cost is an over simplification, the reason arbs trade on transparency information is the expectation of a further price change from liquidity demand that will permit them to unwind the trade at a profit. In this context, proportionality is a reasonable working hypothesis.

²³ Cited in Edelen, Evans and Kadlec [2007] p. 8. In general, trades in “growth” stocks are less costly than trades in “value” stocks, but we ignore this distinction in the table.

²⁴ A recent paper in the market microstructure/execution immediacy literature, Chacko, Jurek and Stafford [2008], analyzes the market impact of S&P 500 additions as an extreme example of the cost of demanding immediate

execution. It will be interesting to compare the results from this approach to results from the comparison of actual transparent executions to the expected cost of an efficient execution suggested in this paper. There is reason to expect complementary results.

²⁵ See Broms and Gastineau [2007].

²⁶ The possible resurrection of sunshine trading to accommodate liquidity demands such as those associated with the apparent de-leveraging of long-short hedge fund portfolios in August 2007 is raised by Khandani and Lo [2007] p. 29. Definitive comparisons of the opaque trading that initiated this market event with the more transparent liquidation of Long-Term Capital Management (LTCM) in 1998 are not yet available. It is likely, however, that the first fund to de-leverage in 2007 fared better than it would have with pre-announcement of its trading intentions – and far better than investors in LTCM fared in 1998.