DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

USING INFORMATION FROM TRADING IN TRADING AND PORTFOLIO MANAGEMENT: TEN YEARS LATER

David J. Leinweber

SOCIAL SCIENCE WORKING PAPER 1135

May 2002

Using Information From Trading In Trading And Portfolio Management: Ten Years Later

DAVID J. LEINWEBER[1]
California Institute of Technology

May 2002 (Revised May 31)

The author is a visiting faculty member in economics at Caltech. At the time the original version of this paper was written, he was Managing Director at First Quadrant [2]*, with responsibility for global active equity portfolios totaling $6 billon. Dr. Leinweber has a long background in electronic trading. He is the inventor of MarketMind, the expert system now incorporated in ITG's Quantex product, and a founder of Codexa, which provided internet information filters for institutional investors and traders. He holds a Ph.D. in applied mathematics from Harvard University and undergraduate degrees from the Massachusetts Institute of Technology.*

The centerpiece of this paper is a comprehensive analysis of all of the trading by a large US pension fund in 1991. The first version was published in the Summer 1995 issue of the Journal of Investing and later incorporated in the AIMR's CFA reading. This updated version, requested by the AIMR, reflects the technological and market structure changes of the last ten years, and adds a new empirical analysis of all the 2001 trading by a $7 billion US equity manager.

Trading technology has changed in many ways, yet many of the same characteristics of institutional trading are seen just as strongly as they were before. The attention given to large difficult orders still shows in lower than expected trading costs. Small "no-brainer" orders still represent the largest component of overall trading costs. These are precisely the type of costs that modern electronic trading systems are designed to reduce.

The idea that trading costs were a substantial drag on performance was a relatively novel in 1991. Today, it is a central concern for many managers, including those profiled here.

Keywords: Transaction costs, trading costs, equity markets, stock markets, electronic order working, liquidity, program trading

[1] djl@caltech.edu, or dleinweber@post.harvard.edu. Author's website: www.hss.caltech.edu/~djl/

[2] www.firstquadrant.com

INTRODUCTION

Trading is the implementation of investment ideas, and the quality of the implementation is as important as the idea itself. For small transactions in large markets, this is not a major concern. As the transaction size grows, or when smaller markets are involved, it becomes important for investment managers and their clients to pay more attention to implementation.

The original version of this article [Leinweber 1995] presented some surprising results from an analysis of a large U.S. pension fund's equity trading, and describes how some of these lessons are being applied in the management of equity portfolios. This update revisits the same ground, after a decade of intense technological change.

REAL PORTFOLIOS AND PAPER PORTFOLIOS

Virtually all equity managers have personal experiences suggesting that paper portfolios outperform real ones. A paper portfolio is an imaginary holding consisting of all the security positions the investor decides to hold, acquired at the midquote price that prevailed at the time they decided to hold them. Paper portfolios incur no commissions, no taxes, no bid-ask spreads, no market impact, and no opportunity costs. Real portfolios incur all of these costs.

The best way to measure transactions costs is to look at the difference between real and paper portfolios. This was first suggested in [Treynor 1981]. Andre Pérold, in a frequently cited paper [Perold 1988], named this difference "the implementation shortfall."

A famous example of the implementation shortfall is the striking difference between the performance of the paper ValueLine portfolio and the real ValueLine fund (see Exhibit 1). From 1979 to 1991, the ValueLine paper portfolio has an annualized return of 26.2%. The real ValueLine fund lags substantially, with an annualized return of only 16.1%. Yet the fund was making the same trades recommended in the newsletter, with only a slight delay after publication (so as not to front-run their mail subscribers). Clearly, something happened. That something is the cost of implementation.

Equity trades and managers are increasingly aware of this and, consequently, increasingly concerned with transaction costs. It is now generally understood that these costs have the potential to erode or eliminate the value added by money managers. Trading is the implementation of investment ideas, and the quality of the implementation is as important as the cleverness of the idea. The growing appreciation of the role of trading in the investment process is reflected in the AIMR's publication of proposed guidelines on trade management and best execution in November 2001.[3]

[3] http://www.aimr.org/pdf/standards/proposed_tmg.pdf

Exhibit 1

Real and Paper Value Line Portfolios

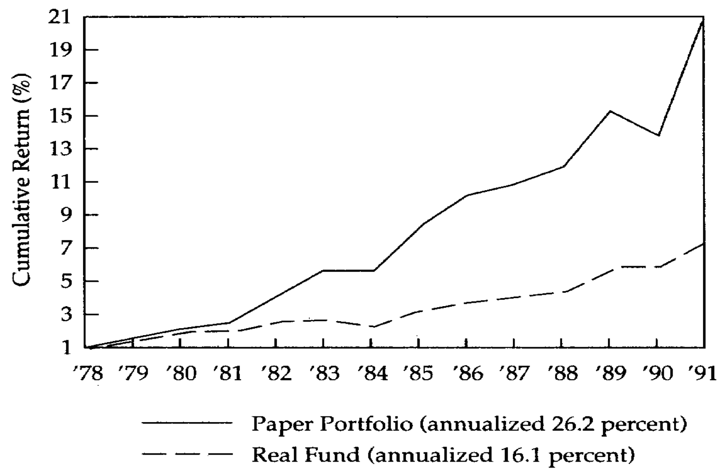

Sources: Value Line; Morningstar; Aronson & Fogler.

TESTING THE UNTESTED WISDOM

There is a large body of untested wisdom about trading. It has not been a particularly popular area of investigation for the academic community (until recently) because of the scarcity of data. In the absence of a well-researched body of empirical and theoretical knowledge, traders have developed rules of thumb and informal guidelines to suggest how to achieve the "best execution." Unfortunately, many of these guidelines are untested or untestable.

When trading was done manually, these guidelines were loosely applied. Today, the growth of electronic trading systems has forced the issue of understanding best execution at a level of detail sufficiently precise to use in an electronic market. People now have to write computer programs to trade well using automated systems. Anyone who has ever written any sort of computer program knows that it is a great method of finding out precisely what you do not know about the subject at hand. Those "minor details" waved off in a casual conversation (or sales pitch) must be filled in before the program will do what it is supposed to do.

Motivated by the desire to use electronic trading systems to achieve best execution for equity transactions, we empirically investigate over 13,000 equity transactions executed by a large U.S. corporate pension fund in 1991. The total value of the transactions is approximately $2 billion.[4]

The analysis is greatly simplified by two facts: All transactions are purchases, and all transactions are completed, so there is no question of how to measure the opportunity costs. Trading cost is simply computed using the implementation shortfall method: the trade price less the decision price plus commissions.

The goals of this study are to empirically test much of the conventional wisdom about trading and to shed light on the question of exactly how we can best use electronic trading systems. Our analysis produces a number of unexpected findings:

- Smaller trades are responsible for a disproportionate share of the costs. The trades expected to be low-cost "no-brainers" are not. (This is observed in both the 1991 and 2001 data.[5])

- Larger trades, generally handled by higher-commission brokers, have lower than expected costs. (Also observed in the 2001 data).

- The expected relationships between management styles and costs are not found.

Other expectations are confirmed:

- Skillful execution reduces costs.

- Patient trading reduces costs.

- Crossing reduces costs.

- Some trades produce transactions profits, which partially offset costs.

TRANSACTION COST PREDICTION

The industry's current practices for transaction cost control are strictly "rearview mirror" approaches. You get a report sometime after you complete a transaction telling you how you did. Getting in front of the problem requires predicting costs before the trade is made. These predictions are needed for individual transactions and for portfolios. Research at MIT [Lo 1993]

---

[4] This updated paper concludes with a brief study of 15,000 transactions in 2001, with a total value exceeding $18 billion, including buys, sells, short sales, and buys to cover. A more in depth analysis of that data is in progress. Check the author's website www.hss.caltech.edu/~djl/ for details.

[5] Omission of an "also observed in the 2001 data" comment doesn't imply the opposite. The analysis was much more cursory.

and [Lo, MacKinlay, and Hausman 1992] is promising in this context. These models can flag problematic trades in advance and provide feedback for portfolio construction decisions.

One trading technique, guaranteed principal bids, incurs costs that are fully known before the trade is made. The technique itself, in effect, embodies a perfect “forecast” of trading costs.

Reliable transaction cost forecasts can be applied in two places. The first is in the often discussed but seldom-observed integration of portfolio management and trading. Most portfolio construction tools and optimizers use overly simple assumptions about transaction costs, e.g., they are the same for all stocks and are expressed as so many cents per share or a fixed percentage of the order size. These assumptions are clearly at variance with the real world. The portfolios that would emerge from these systems if they incorporated more realistic estimates of transaction costs are very different from the portfolios produced under the naïve assumptions.

The cost history of previous transactions in an investment process can be used to explicitly predict costs or otherwise influence trading methods used to implement that process in the future. This detailed and explicit consideration of transaction costs is an important link between trading and portfolio management.

Costs can be measured after the fact, and used to plan future trading. They can also be measured as a trade progresses and used to determine how it will be completed. This dynamic measurement of costs and feedback into a trading program is a valuable cost control technique available only to traders using electronic systems capable of keeping up with the process in real time. Measurement is the first step to control on multiple time scales. In 1991, transaction cost measurement was truly a boutique service. Ten years later, there are many alternatives for managers looking to do this – brokers, consultants, and internal monitoring systems provide a much more detailed and timely means of feedback for cost control.

HEAT, LIGHT, AND TRANSACTION COSTS

This article began with a discussion of trading costs illustrated by comparing real and paper portfolios. The implementation shortfall method described in that context is the cost measure here, but it is not the only measure that can be used to measure costs. The whole subject of transaction cost measurement tends to generate quite a bit more heat than light. Many measures have been proposed. The simplest of these are really proxies for transaction costs. Examples include bid-ask spread, numbers of trades occurring within the bid-ask spread and the like. These are useful in comparing entire markets (e.g. [Chan 1997], [Bessembinder 1997], and [Keim and Madhavan 1998]), but they are less informative in telling us about the transactions costs incurred by a particular investor. A bid-ask spread of a penny is nice, but largely irrelevant to when the size associated with that spread is 200 shares, and the trade in question is for 2 million shares.

True transaction costs are fundamentally immeasurable (see [Wagner 1990]). This is because they are the difference between the price you paid and the price that would have prevailed if you *had not transacted*. We can never observe this price, so we can never measure true cost. If this were the last word on the subject, we could all go home now, but it is not.

The implementation shortfall method has been widely accepted by both academic economists and market practitioners as a good surrogate measure for true (and therefore unobservable) transaction costs. It goes beyond quote data, which are the same for all investors, to consider the circumstances applicable to the actual trade made by a particular investor. The key feature of the implementation shortfall is that it is based on difference between the price the investor decided to transact, and the total price of the transaction, so it fully captures market impact, commissions, taxes and all other costs of the completed trade.

It is occasionally said that it is futile to attempt to control transaction costs. You cannot measure them precisely, so why try? This is specious. The ability to make precise measurements of a phenomenon is not required to influence or control it.

Recall the Heisenberg uncertainty principle, which states that it is fundamentally impossible to perfectly (and simultaneously) measure the position and velocity of a particle such as an election. A single glance at a television or computer screen confirms that we can indeed control the little devils well enough to broadcast a Dodgers game or graph the S&P 500, tick by tick, on a computer screen in 16 million colors, even if we cannot measure the electrons' position and velocity with perfect precision. The analogy to transaction costs: Even if we cannot measure the costs *precisely*, we can influence them by the actions we take in executing trades.

An excellent current update to the literature and technology of transaction cost controls is found in [Bruce 2002].

INFORMATION TECHNOLOGY AND MARKETS

The 1975 legislative mandate to develop a U.S. national market system encouraged "maximum reliance on computer and communications technologies." This has happened with a vengeance, both in the U.S. and around the world. In 1991, the NYSE's DOT system handled two-thirds of the orders on the world's largest stock exchange. By 2001, this proportion had risen to over 90% of the orders. NASDAQ used automation to transform itself from a sleepy market for small companies to the second largest U.S. equity market.

At this point, the previous version of this paper listed a large number of then-current electronic execution systems. Some are thriving (e.g. POSIT), some are gone (e.g. the Arizona Stock Exchange and CompBid). One of the largest efforts to create an electronic market, Optimark, had not even begun in 1991. In the intervening years, it attracted a prodigious amount of high profile venture capital and publicity, and then disappeared. Others have taken their place. The market for electronic equity executions has been a prime example of Schumpeter's view of capitalism as "a process of creative destruction."

ECNs[6], represented only by Instinet in 1991, have proliferated and flourished. Collectively, ECNs now account for approximately a third of the volume in NASDAQ stocks. The repeal of the NYSE's Rule 390 has seen the beginning of trading in listed stocks on ECNs.

The NYSE is now the world's only major equity market with a traditional trading floor. However they have greatly expanded their automated systems, providing both direct execution services, and access to floor liquidity via electronic channels. Multiple electronic channels are now used for 90% of NYSE orders, representing over 50% of the volume.

The NYSE limit order book, which could be seen only on the floor for over 200 years, was made visible to traders using the OpenBook system, in 2002. This change followed the move to decimal pricing, which had reduced the amount of liquidity information conveyed by the inside best bid and ask prices.

Interestingly, as technology allows the dissemination of limit orders, it also allows traders to reduce their use of these orders. Keeping one's cards hidden is a central feature of all trading strategies, but one also needs to be present in the market to capture liquidity. In the past, this tradeoff encouraged the placement of limit orders. Now with the increasing ease of automated market monitoring and electronic order entry, it is very feasible to keep your limit orders to yourself. This is often referred to as "hidden liquidity".

With the clarity of hindsight we can see some patterns. Price discovery has been a hard sell. AZX and Optimark both sought to help discover prices, but failed to attract enough volume to fuel that process. Primary markets, NASDAQ and the NYSE, provide price discovery as a free good to all market participants[7]. Anonymity is good. Institutional trading is a game of not showing one's hand. Simplicity wins over complexity.

The ubiquity of the Internet and its universal standards for high-speed communications involving multiple machines has fostered the rapid growth of ECNs and the NYSE/SIAC family electronic trading systems for listed stocks. Parallel developments are evident throughout the world.

A new generation of even more sophisticated real-time trading systems are capable of implementing trading strategies to be employed across these electronic markets. Current state of the art ideas are embodied in systems like (but not limited to) Quantex, Flextrade, LavaTrade and UNX.

---

[6] "Electronic Communication Networks" – A very generic name for systems designed to match up natural buyers and sellers. Examples are Archipelago, Brut, Redibook, and Island. In 2002, Archipelago, one of the largest ECNs, legally became a registered stock exchange.

[7] The primary markets are acutely aware of this "free rider" issue. They are eager to find ways to derive additional revenue what has become virtually cost-free resource to their competitors.

Outside the U.S., we see even more dramatic evidence of the technological transformation of markets. The trading floor of the London exchange is abandoned, rented out, and replaced by information systems. Toronto's trading floor closed in 1995, and Tokyo's in 1998. Both replaced by a fully electronic trading systems.

Guaranteed principal bid portfolio trading is made possible by technology. U.S. and international investors seeking to control costs increasingly use it. Detailed analysis of portfolios, given only their characteristics and the balancing of risk across a large number of portfolios, would be impossible without modern computational tools.

PUTTING PROGRESS IN TECHNOLOGY IN PERSPECTIVE

Advancing technology makes all of this real. In the U.S., the consolidated trade and quote feeds and the Intermarket Trading System couple diverse market segments. By the early 90s, well-wired market participants received machine-readable feeds covering hundreds of world markets at rates up to fifty-six kilobits a second. Even though that seems quaint in 2002, it corresponds to a full transmission of the entire *Wall Street Journal* financial listings every three seconds. In 2001, they do this on hand held wireless devices. Landlines are ten or a hundred times faster. The demands of market data, a text stream, are a negligible load on the bandwidth of current networks.

What can anyone possibly do with information arriving in such quantities and speeds on these "fire hose" data feeds? The obvious answer is that *people* alone cannot do anything with that much information. The receiving end always involves a computer. The kinds of tasks we can do with computers have grown so incredibly that, to use them effectively, we have to change the way we think about them.

No technology in history has progressed as rapidly as electronic computing. Let's look at this in more detail[8]. Two reasonable ways to measure the capacity of a computer are by its memory size (in megabytes) and its speed in MIPS (million of instructions per second, in the ads, or meaningless index of processor speed, in the EE department). A reasonable way to measure its cost is by its price (in current dollars). Thus, a good figure of merit for computers is (Memory x Speed)/Price.

A 1960s-era PDP-1 computer filled a large room. It operated at about 0.1 MIPS, had 0.008 megabytes of memory (which the salesman would have called 8K), and cost $250,000. Our cost/performance measure turns out to be 0.0000000032. A modern high-end workstation (the kind you find by the truckload on Wall Street) runs at 100 MIPS, has 256 megabytes of memory, fits in a desk drawer, and costs about $20,000. Our figure of merit is now 1.28. This is

---

[8] Gordon Moore, a founder of Intel Corp, popularized this analogy. My apologies to readers who have heard this before. It bears repeating.

an improvement from the early 1960s by a factor of 400 million (which readers of this Journal might call 4,000,000,000,000 basis points). Exhibit 2 shows this growth in a fairly standard way.

To appreciate what this really means, consider what would happen if the same rate of progress had occurred in the automobile industry during the same period. Let's define a similar figure of merit for two cars as (Mileage x Cruising Speed)/Cost. Two good things multiplied together and divided by a bad thing.

Exhibit 2

Computing Power Per Dollar

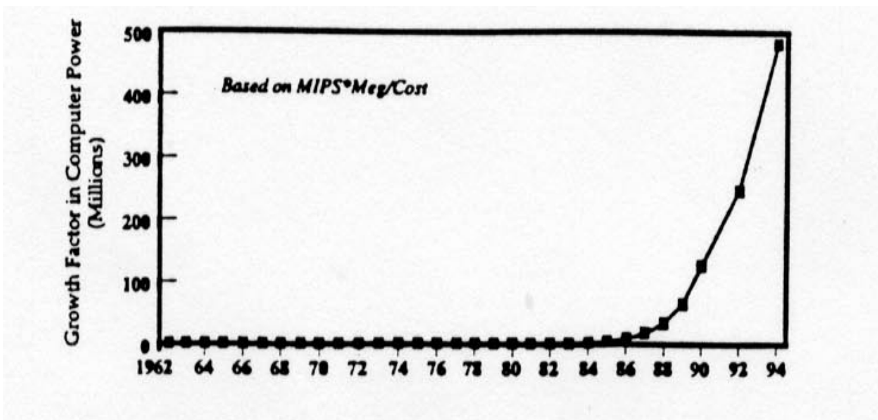

Note: Based on speed × memory/cost.

We start with a typical 1960s Chevrolet. It gets 10 mpg, cruises at 75 mpg, and costs $2,000. There are an infinite number of ways to scale up the car's figure of merit by a factor of 400 million to match the computer's progress. Here is one: the 1994 scaled-up car gets 2,000 miles per gallon, cruises at 7,500 mph, and costs twelve cents. This is not really what we would consider a car at all; it is a virtually costless, disposable land rocket. We would use this in very different ways from how we use today's cars. It would totally change our lives, our economy, and our investments. (Of course, we can keep updating these numbers, applying the Moore's Law factor of two every eighteen months. In 2001, that gives us another factor of ten. The twelve-cent car now costs 1.2 cents, or it gets 20,000 miles per gallon. This will continue for at least another ten or twenty years, and probably longer. So it becomes a tedious exercise. You all get the idea here.)

This factor of hundreds of millions typifies the magnitude of the challenge facing the securities industry regarding computers: *If you had everything computationally, where would you put it financially?* One place to put it is in the improvement of the trading process, and a reasonable question is: "How?" This is the motivation for our attempt to understand the nature of the transaction process better and use that factor of 400 million to reduce the cost of transacting. Here are the details.

A $2 Billion Experiment

We examine 13,651 equity purchase transactions, totaling nearly $2 billion, made by one of the largest U.S. corporate pension plans in 1991. Plexus Group[9], a firm that analyzed transactions costs, with the cooperation of the fund manager, provided the data. Trade sizes range from 100 shares to blocks of more than 400,000. Both active and passive management styles are represented. All orders in this sample are filled, some on the first day after the decision to trade, some up to twenty-one trading days later. Because all orders are eventually filled, no opportunity costs are observed here. This simplifies the measurement and interpretation of costs, as there is no need to define opportunity costs.

We set out to test the validity of the conventional wisdom in several areas: specifically, the relationships between transaction costs and 1) trade size relative to market capitalization, 2) trade size relative to average trading volume, 3) management style, 4) patience in trading, and 5) use of crossing networks.

Some findings are surprising, although they must be taken with the appropriate caveat: These transactions all come from a single pension fund, and generalization to all funds may be premature. However, as these results have become more widely known, several other individuals, from the United States and Japan, have told me they have observed very similar effects in their independent analyses of trading data. They have sent in charts of thousands of their trades that look remarkably like the charts in this article, so we believe it is safe to say these are not completely unique observations.

At the end of this updated paper we examine (in less detail) an even larger and more varied set of US Equity transactions, with a total market value exceeding $18 billion.

Costs and Trade Size

What should we expect regarding costs and block size? The most reasonable sounding model is first described in [Loeb 1983] and extended in [Loeb 1991]. This very plausible model is based on queries to brokers and market makers for blocks valued between $5,000 and $20 million at various levels of market capitalization. The 1991 paper generalized the earlier result by fitting an equation to the data.

---

[9] www.plexusgroup.com

The Loeb function predicts transaction cost (in percentage of value) as a function of trade size expressed as a percentage of shares outstanding and market capitalization. Exhibit 3 shows these predictions. Our expectation is that a scatter plot of cost versus size as a percentage of shares outstanding for the 13,651 transactions in our sample would lie along these lines.

Exhibit 3

Expectations – Costs versus Block Size

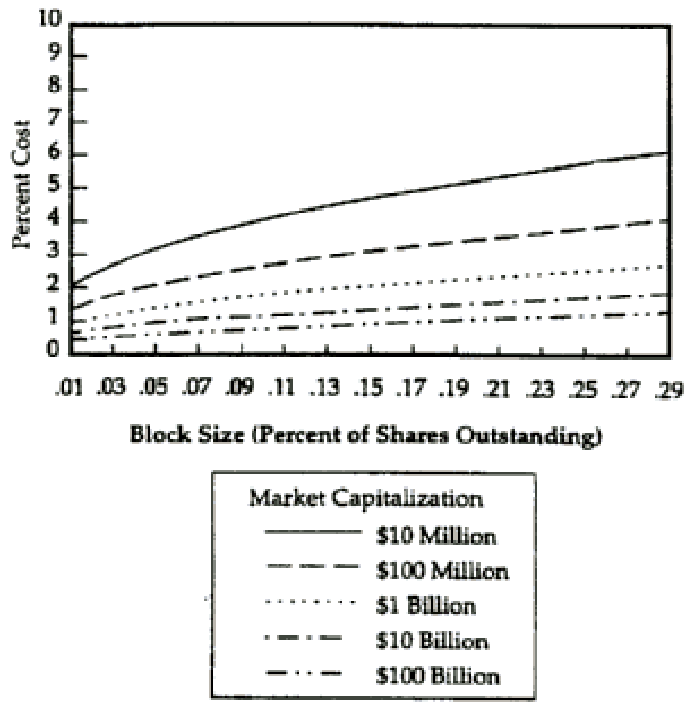

Exhibit 4 shows a pattern clearly different from our expectations. The largest trades have costs much lower than those predicted by the model. The highest and lowest percentage costs are associated with the smaller trades. Many trades, of all sizes, have negative costs; that is, they produce transaction profits. Overall, costs rise slightly as trade size increases, but they rise less than expected. What is responsible for this? It may be that these are largely actual trades, not suppositions. Apparently, when real orders are filled using real money, block traders can deliver better performance than the Loeb model predicts. These lower-than-expected costs for the largest orders seem to reflect the self-selection resulting from the attention traders give them, they are more likely to take place when they can be done most effectively.

Exhibit 4

Observations – Costs Versus Block Size

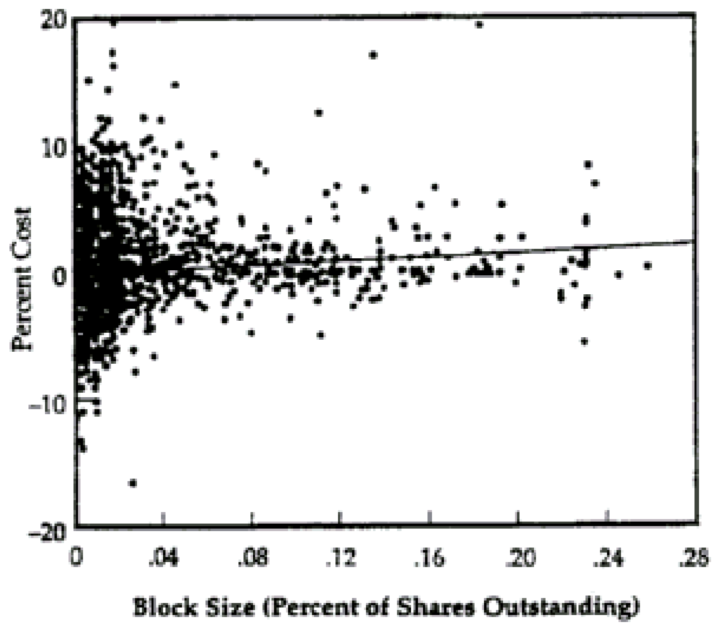

One objection to the use of percentage of shares outstanding as a standardized measure of order size is that it is a poor indicator of the actual difficulty of a trade. Many shares are closely held and do not trade very often. Looking at trade size as a percentage of average daily volume may be more appropriate. A scatter plot of the same transactions with the size as a percentage of the three-day average volume (prior to the decision date) is shown in Exhibit 5. This appears to be a better measure of size, but it still fails to explain the surprising bulge in both trading costs and profits for the smaller transactions. For these smaller trades, the costs substantially offset the profits, suggesting that the cumulative price of "no-brainer" executions is much higher than realized.

Exhibit 5

Observations – Costs Versus Block Size

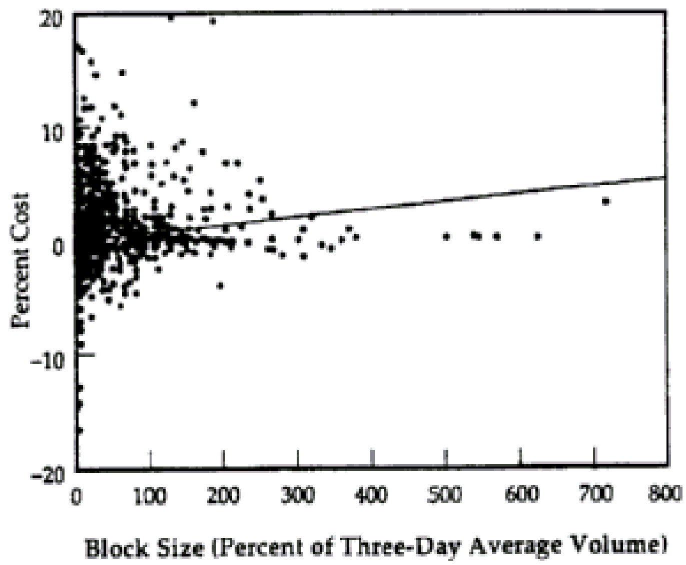

This is illustrated when we look at net percentage trading cost by order size, as shown in Exhibit 6. What does this mean in terms of actual dollar cost to the pension fund? The answer is seen in Exhibit 7, which illustrates the total transaction costs by trade size. The small orders represent the largest single contribution to total trading costs.

Exhibit 6

Net Trading Loss by Order Size

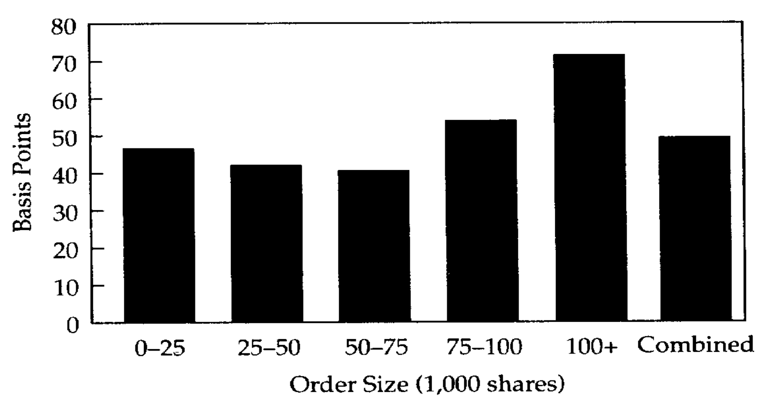

Exhibit 7

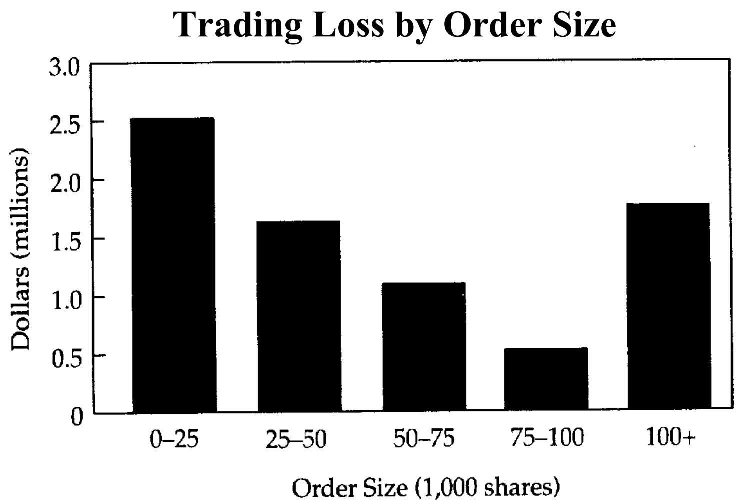

Costs and Management Style

What do we expect to observe when we look at transaction costs separated by management style? An excellent discussion of this subject is found in [Collins and Fabozzi 1991]. Patient disciplines, such as value and growth investing, have longer time horizons and are expected to have lower transaction costs. Earnings strategies depend on quicker execution to capture the market's reaction to differences between expected and actual earnings. Index funds tracking small-capitalization stocks would be expected to have larger costs because of the characteristics of the smaller stocks compromising those indexes. Execution costs for high-capitalization index funds, such as those tracking the S&P 500, are expected to depend on the funds' use of index futures to manage costs.

These expectations are summarized in Exhibit 8. Observations of the actual transactions for different management styles are shown in Exhibit 9. Management styles expected to exhibit the highest costs have the lowest costs, and vice versa. This surprising result suggests that something may be peculiar to this particular fund (as suggested by the earlier caveats), or that our understanding of the relationship between management style and costs is in need of refinement, or that fund managers may not be rigorously adhering to their target style.

Exhibit 8

Expectations — Costs and Management Style

Management Style	Trade Motivation	Liquidity Demands	Execution Cost	Opportunity Cost
Value	Value	Low	Low	Low
Growth	Value	Low	Low	Low
Earnings Surprise	Information	High	High	High
Index Fund, Large-Cap[a]	Passive	Variable	Variable	High
Index Fund, Small-Cap	Passive	High	High	High

[a]The costs associated with some investment strategies that use futures can be low, despite high opportunity costs.

Exhibit 9

Observations – Costs versus Management Style

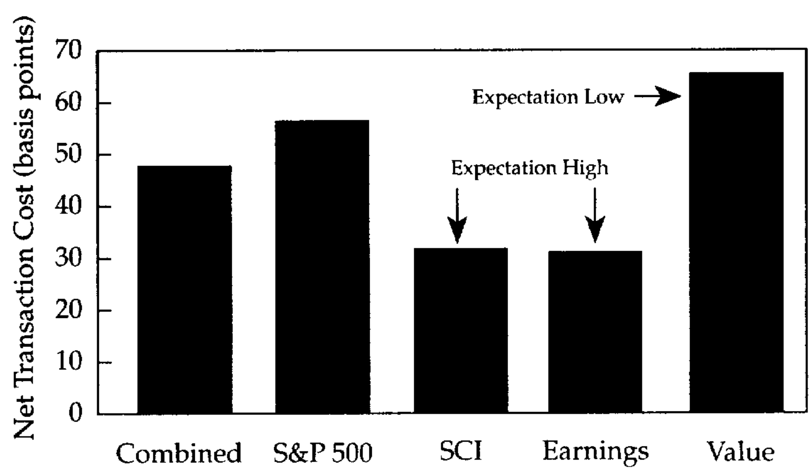

COSTS AND PATIENCE

Our expectation is that patience will be rewarded. Patient trading allows short-term market volatility to work in your favor. You have time to exploit crossing systems, which typically transact only once or twice a day. The market impact of patient trades should be less than large trades executed in a hurry.

The reasons underlying our expectations regarding patience are illustrated in Exhibit 10, taken from [Collins and Fabozzi 1991]. The figure shows how opportunity cost and execution cost add up to the total cost of transacting (or not transacting). Recall, however, that all the transactions we study are completed, so no opportunity cost is involved and we would expect to observe only the execution cost curve.

Exhibit 10

Expectations – Costs versus Time

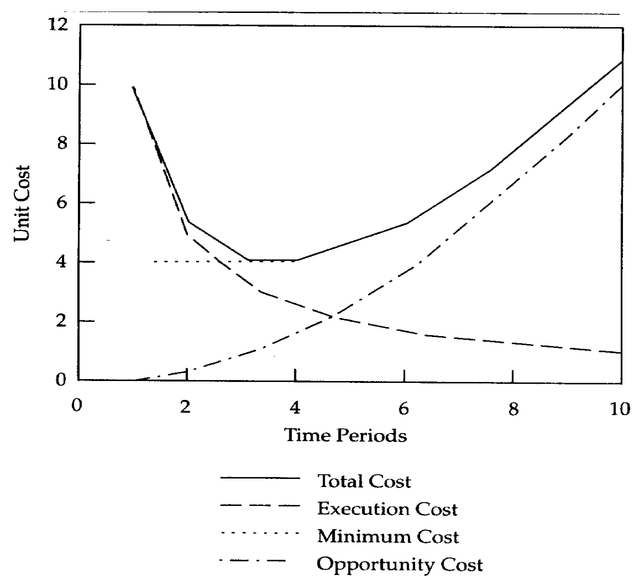

Source: Bruce Collins and Frank Fabozzi, "A Methodology for Measuring Transaction Costs," *Financial Analysts Journal* (March/April 1991):27–36.

These expectations are reinforced by the results of a real exercise inpatient trading reported in [Bodurtha and Quinn 1990], as shown in Exhibit 11. This is an actual series of trades to construct a 350-stock portfolio to track the Russell 2000 Index. The returns on the real and paper portfolios are shown in the upper chart, along with the implementation shortfall (the difference between these two curves). This shortfall is almost zero for the first fifty days of patient trading. During the period when the trading was done patiently, the real portfolio tracks the paper, with very low transaction costs. When patience ran thin, the implementation shortfall rose dramatically, clear evidence that patient trading can pay.

Exhibit 11

Results of Patient Program Trading

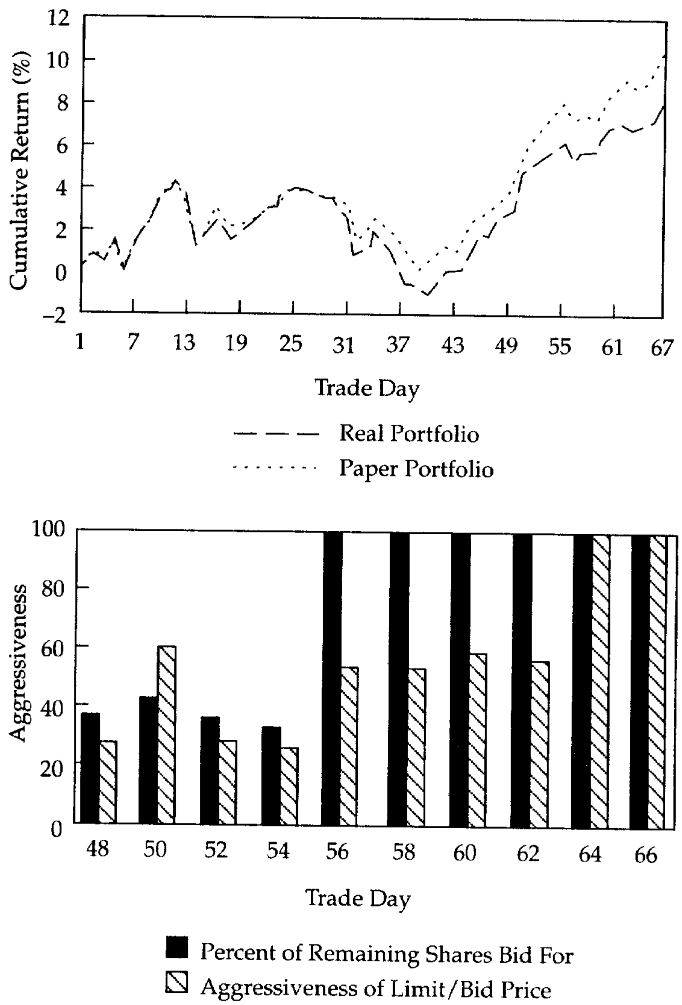

Source: Stephen Bodurtha and Thomas Quinn, "Does Patient Program Trading Really Pay?" *Financial Analysts Journal* (May/June 1990):35–42.

The lower chart shows that, after fifty days, the traders were losing patience and became more aggressive on size, increasing the percentage of remaining shares bid for from roughly 35% to 100% of the total remaining, and roughly doubling this "aggressiveness" on the price of their limit orders[10]. At day 63, they really wanted to get this over with, and again increased their aggressiveness on price.

This set of transactions requires an extraordinary amount of patience, sixty-seven trading days. Is this necessary? Probably not, and the reason is an excellent example of how to use information from trading in portfolio management. An incredibly large number of 350 stock portfolios will track the Russell 2000, so there is no reason to pick a fixed set of 350 names and stick with it for sixty-seven days. They operate in the traditional manner: a portfolio manager constructs the list, passes it on to the trader, and waits for completion.

Exhibit 12 is a look at the actual trades, typically quite large, which were spread over multiple days. It shows that patience was rewarded.

Exhibit 12

Mean Cost versus Days Between Decision &Transaction

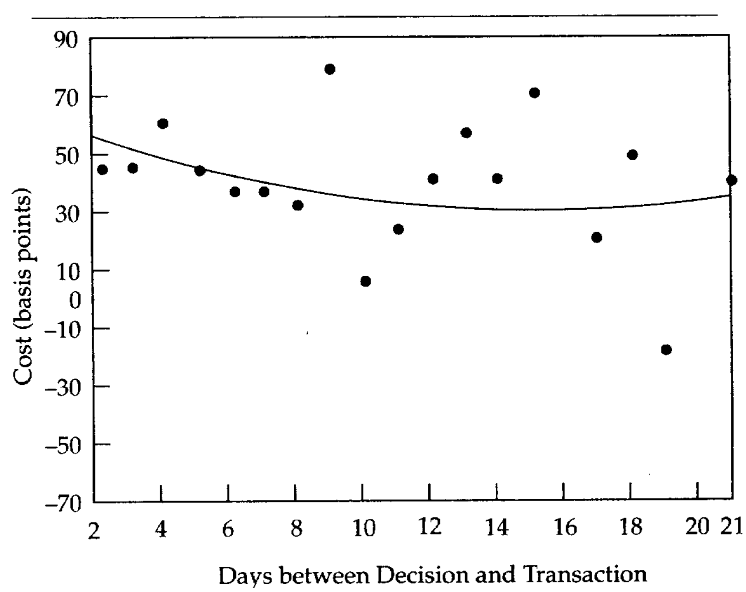

A somewhat exotic alternative is to work the process as a substitution order.[11] To do this, a larger set of names, perhaps 1,000, would have been generated in the first place. The basic

[10] Aggressiveness on price is measured by the position of the limit order price within the quote prevailing at the time the order is submitted. Aggressiveness on size is measured by the percentage of shares remaining to buy, represented by the limit order size.

[11] This was first suggested (and implemented) by Evan Schulman at Batterymarch in the '70s, and is described in the Harvard Business School Case 1-286-113

patient trading tactic would be the same, but each day's fills are fed back into the portfolio construction system as current holdings to generate an updated trade list. Any of the "substitutable" names can be used, rather than waiting for low-cost opportunities to trade the original fixed list. Undoubtedly, the process could have been completed in less time, with comparable lower costs.

Why is this not done more often? Many investors do not want to substitute one stock for another. This is feasible only for certain passive portfolios, or active managers looking to adjust industry or sector exposures. Even in 2001, most standard trading systems do not support it, but specialized automated systems can and do allow management of thousands of dynamically changing orders, modifying or canceling them as the market moves and as executions occur. It is time to fire up the twelve cent Ferrari, i.e., the computerized trading/portfolio management support tolls made possible by that factor of 400 million in technological growth. Today's automated systems can, and do, manage the task of controlling this type of electronic order working to achieve best execution in a broader sense.

COSTS AND CROSSING

One reason patience pays is that patient traders participate in crossing networks more often, and crossing should reduce execution costs. The reason for this expectation is illustrated in Exhibit 13, which shows the components of trading costs. The two major crossing networks are ITG's POSIT, which does multiple intraday, midquote crosses, and Instinet's Global Crossing Network, which does an after-the-close, closing price cross. Both systems have low commissions because the intermediaries are removed from the trading process. No price concession for size in excess of the quoted size exists, because no size is quoted. The price is independent of transaction size. Both have grown in volume and scope since the first version of this paper appeared. International versions of both systems are now operating and gaining traction

Exhibit 13

Expectations – Crossing and Transaction Costs

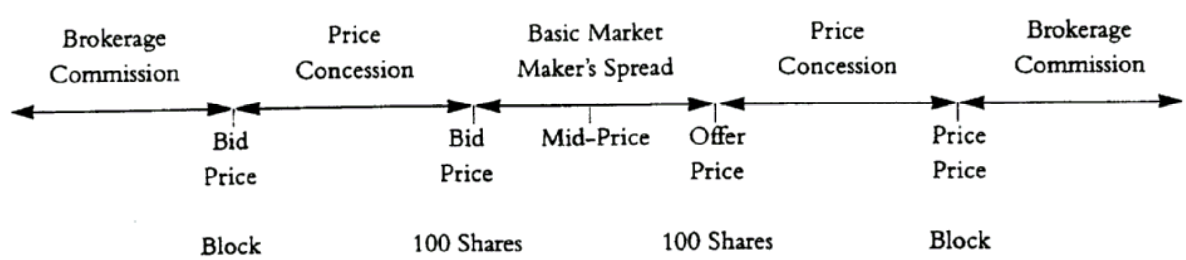

Factors lowering costs for crossed trades include mid-quote or closing price, no concession, and low commissions. (Source: Loeb 1983)

RECAPPING THE PROBLEM AND FINDING A SOLUTION

When we examine a large number of actual transactions made by a large U.S. pension fund, we find:

- The implementation shortfall costs for large trades are much lower than expected.
- Many small to midsize trades generate a transaction profit.
- A large number of unexpectedly costly small to mid-sized trades make a disproportionate contribution to the total implementation shortfall.
- New trading technology is designed to help investment firms control the costs of the smaller trades which represent the largest contribution to their total transaction costs.

If the larger trades are being executed well, with costs below reasonable expectations, and some trades have negative implementation costs (i.e., profits), then the problem illustrated in Exhibit 14 is fairly obvious. The solution to this problem is to find those thousands of costly transactions (in the upper left portion of the figure), and those trading tactics, market channels, and brokers have been most effective in executing similar trades in similar circumstances, with similar motivations.

Exhibit 14

Problem Trades

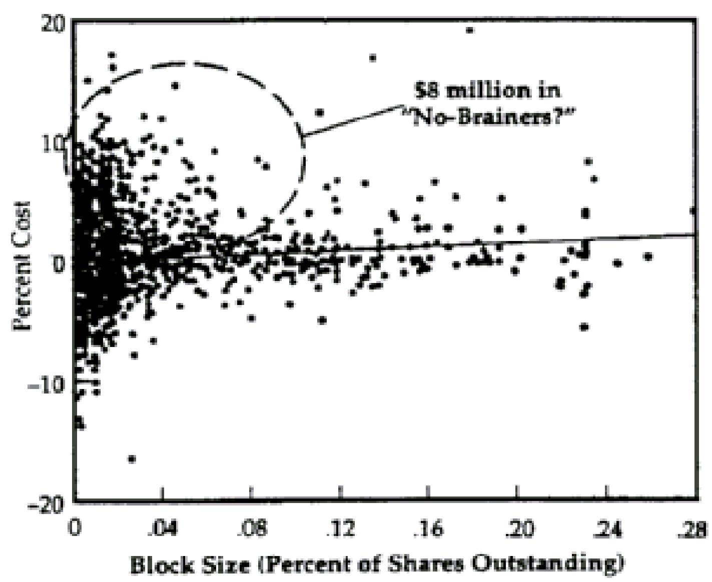

COST REDUCTION LESSONS

This analysis suggests several ways to reduce costs:

- *Value is added by skilled equity block traders.* Because of constraints on their time and availability, these traders can handle only a small fraction of the total order flow. If the high-cost smaller trades, which would not normally be sent to block desks, could be given the benefit of their attention, we could expect better execution.

- *Patient trading and automated liquidity-seeking trading strategies may be appropriate to keep the transaction costs below the anticipated alpha added by the manager.* Electronic trading and order working systems allow the use of computerized trading strategies that would be impossible by other means. Price-sensitive demand schedules and substitutions are also useful here.

Complex trading techniques incorporating multiple layered strategies allow a variety of techniques to be used together, choosing which ones to apply based on market conditions and cost feedback information, is described in [Leinweber 1994]. The basic idea behind these electronic order-working strategies is to use information from current and prior trading to directly affect the outcome of the trading process. Electronic trading systems provide a uniquely powerful means of capturing this information and exploiting it in ongoing trading. Electronic order working is a general term for this type of trading. The basic ideas behind electronic order working are:

- Exploit the multiple execution channels available today.
- Allow short-term market volatility to work in your favor.
- Apply the techniques simultaneously to a large number of orders.
- Incorporate feedback from the results of trading strategies on multiple time scales to refine the performance of those strategies. As stated (and elaborated in great detail) in [Grossman 1989], the "very activity of trading produces information which affects its outcome."

- *Identification of expensive problem trades before they are made.* There are two aspects to this. A general approach is to avoid trading in illiquid stocks by careful consideration of the investable universe and the sizes one is willing to trade within that universe. A more particular approach is to try to predict costs for each transaction, and use this information in portfolio decisions.

Average liquidity is relatively simple to estimate. Short term variations are harder to predict accurately. However the ability to use probe trades, which measure liquidity in real-time, and to quickly incorporate the results of recent trades, brings a timely source of data to the process.

GUARANTEED PRINCIPAL BIDS

An important cost control strategy is guaranteed principal bid (or package) trading. GPB traders receive a firm bid on an entire portfolio from a broker based on disclosure of *characteristics* of the portfolio, not the disclosure of its constituents. Typically, this transaction is priced in advance, in cents per share or basis points off the closing prices for the stocks. This can be regarded by the portfolio manager as a perfectly accurate prediction of transaction costs.

Brokers seek to match a large number of such portfolios, and thus reduce the risk that their cost of trading the unmatched stocks in the market will exceed their revenue from the guaranteed principal bid. One of the difficulties of GPB trading from the point of view of the manager is the difficulty of obtaining bids from a variety of brokers who all want different information describing the portfolio. CompBid, an electronic market to obtain multiple bids on a portfolio trade, operated for a time in the nineties, but has vanished. Certain agency-only brokers now provide a similar service.

COST PREDICTION AND GUARANTEED PRINCIPAL BIDS

If we could predict trading costs with perfect accuracy, we would just add them to our decision prices and make perfectly informed decisions. There is some evidence of predictability, but these predictions are far from perfect. When we use a guaranteed principal bid, we can predict costs with perfect accuracy. That is the guarantee.

This is illustrated in Exhibit 15. The solid line shows the expected transaction costs. The scatter points show the actual costs. This time there are no surprises. This is particularly valuable for quantitatively managed portfolios, which can be selected with known costs going into the optimizer or other system used in the portfolio construction process.

Many managers rely heavily on GPBs. This trading technique is not a magic bullet. It entails a close relationship with the brokers to define the liquid universe and portfolio characteristics that result in acceptable costs. This is particularly true for market-neutral funds, as there is far less liquidity on the short side of the market in many countries.

Exhibit 15

Guaranteed Bids Remove Uncertainty

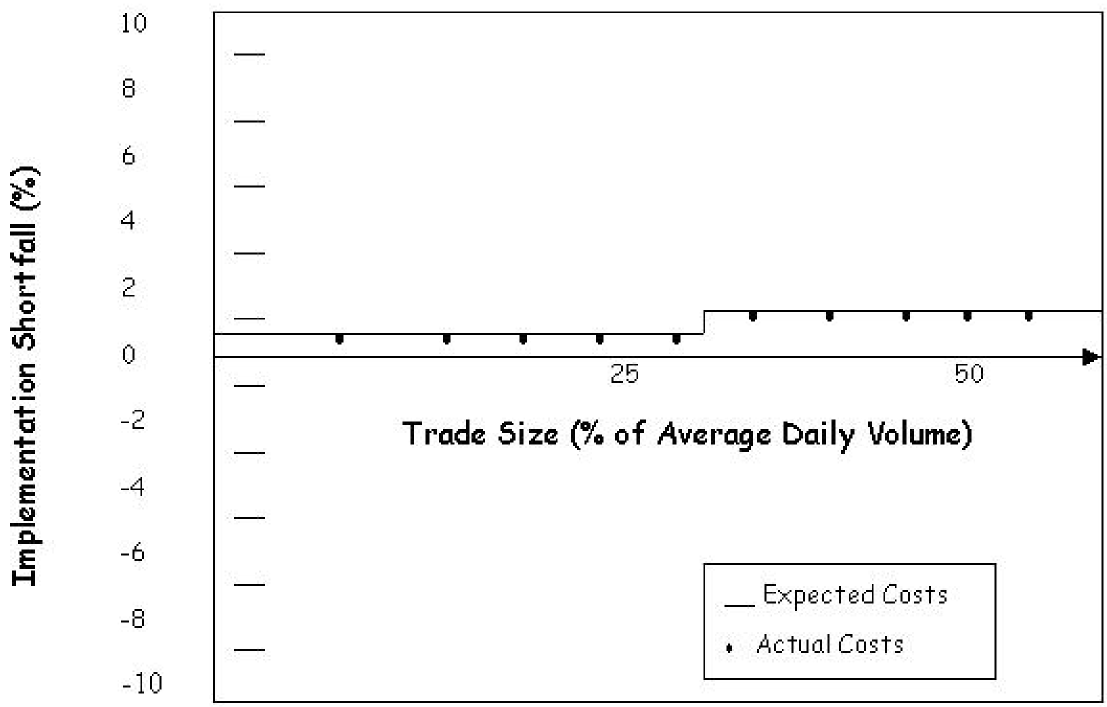

By carefully structuring GPB trades, transaction costs for these portfolios have been held to less than half the reported averages for institutional equity portfolios in their respective countries (see [Pérold and Sirri 1996]).

Are GPBs a free lunch for equity managers? No. The guarantee typically applies to only one portfolio at a time. While there will be no unanticipated market impact for each individual portfolio, brokers are under no obligation to maintain the same pricing for subsequent portfolio trades. The impact costs can be just as real, but delayed. New electronic trading systems increase the liquidity and efficiency of the GPB market by lowering barriers to entry and increasing the likelihood that dealers will find matching transactions.[7]

CHANNELS AND TOOLS FOR ELECTRONIC EXECUTIONS

We have discussed using electronic trading systems as a means to reduce costs. Traders face a growing, overlapping, and often bewildering array of choices in this regard. Chris Keith, formerly of the NYSE, describes the situation as "liquidity divided by confusion equals a constant." Buy-side traders are increasingly motivated to disintermediate their trades, at the same time that the brokers who have been the intermediaries are trying to attract more of them. They face shifting and expanding array of diverse execution alternatives, and technologies to manage them. A thorough review of these alternatives would triple the size of the article, but it is worthwhile to include a perfunctory scan of some of the systems being used for electronic trading of U.S. equities today.

- Flextrade[12] includes a variety of analytic tools and access to multiple execution channels, including broker desks. There are preconfigured and customizable order-working programs for a range of trading situations. The technology can also be used for European and Japanese equities.

- ITG's Quantex[13] system is a versatile means to access execution paths and to direct executions using market data and feedback from trades in process. Many electronic order-working strategies can be implemented. Quantex traders can choose to execute orders electronically, in the POSIT cross, or manually through brokerage desks. ITG also operates in Australia, Europe and Canada.

- Instinet's Order Management System[14] provides access to both the continuous Instinet market and the crossing networks, which cross orders after market hours at closing prices. The Order Management System can work orders using a variety of intraday strategies, with residuals being sent to the cross. International market access is supported.

- Lava Trading[15] systems consolidate market data information from multiple ECNs and exchanges into a single, high-performance data feed. Traders can access this information and submit orders using a unified order execution interface.

- Universal Network Exchange[16], UNX is a modern ASP (Internet Application Service Provider) system to provide sophisticated, customizable tools for traders to operate effectively and anonymously across a wide range of liquidity sources for US equities. Direct, VPN and web access are supported. The core of the system is a sophisticated algorithm that sweeps multiple execution venues to discover hidden liquidity within traders' constraints.

---

[12] http://www.flextrade.com

[13] http://www.itginc.com/products/quantex/quantex.html. For more details, see Leinweber and Beinart [1995].

[14] http://www.instinet.com/equity_marketplace/oms/overview.shtml

[15] http://www.lavatrading.com

[16] http://www.unx.com

This is not an exhaustive list of the choices for electronic equity executions and does not begin to cover the alternatives in the depth they deserve. These systems differ in the type of order working control they provide and, in many cases, in the structure of the combined markets they create. The important point is that liquidity is where you find it, and there are many places and ways to look for it.

TEN YEARS AFTER: A LARGE SCALE LOOK AT INSTITUTIONAL TRADING IN 2001

Ten years ago, when we first saw the trumpet shaped scatter plots of implementation shortfall versus order size (exhibits 4 and 5) our first reaction was that we must have made an awful mistake. Surely, the largest costs wouldn't be associated with the smallest trades. So we cross checked everything we had three different ways, going back to the original transaction records to make sure what we were seeing was really there. It all matched to the penny, they were real.

When the AIMR decided to include the first version of "Using Information from Trading…" in CFA readings, it resulted in a longer and wider circulation than most articles. Consequently, quite a few people were inspired to make their own scatter plots of transaction costs and order size[17]. They are remarkably similar. Over the years, I've received nice collection of trumpet shaped diagrams, mostly from the US, but Canada, Europe, Australia and Japan are there too.[18]

A new dataset of over 14,000 US equity trades, totaling a healthy $18 billon, was made available just in time to be used for the anniversary revision of this paper. In many ways, it is much richer than the original – it includes all 2001 buys, sells, short-sales and buys-to-cover by Aronson+Partners[19], a multi-strategy US institutional manager. The picture is there in Exhibit 16, the same trumpet pattern appears again.

The data summarized in this one picture is certainly worthy of a more in-depth analysis than time and space allow here[20]. But a slightly closer look is informative. Trade sizes averaged 26% of a day's volume, ranging up to 841%. Costs for individual trades vary widely, ranging

---

[17] If you are inclined to do this, don't bother with Excel or PowerPoint, unless you have only 255 trades. You'll need some kind of industrial strength statistical package to handle thousands of points. Try Matlab, SAS, or S-plus.

[18] The last ten years have also seen a greater focus on global equity transaction costs. See [Domowitz, Glen and Madhavan 2001] and [Perold and Sirri 1997]

[19] www.aronsonpartners.com

[20] This is in fact, a work in progress, and will be forthcoming later in 2002. Watch www.hss.caltech.edu/~djl/

from a high of 76% to a 42% profit (these, and a few other extreme outliers are not shown on the chart).

Exhibit 16

$18 Billion in U.S. Equity Trades (2001)

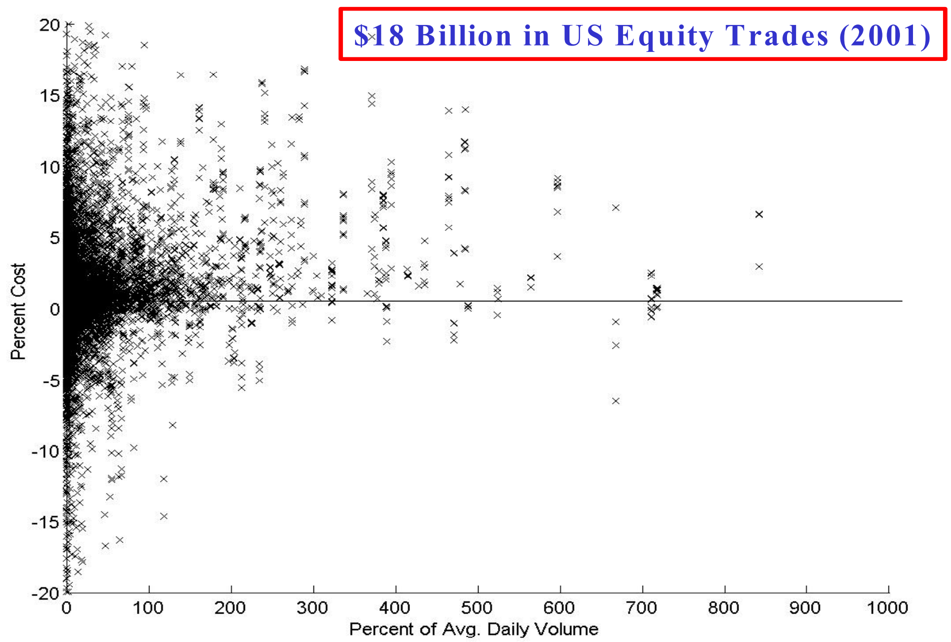

Source: Aronson+Partners

The outliers are unimportant. The aggregate message here is that there are implementation shortfall trading costs totaling $15.1 million. This is offset by trading profits of $8.8 million, for a net trading cost of $6.3 million for the year, an average of 35 basis points.

$6.1 million in transactions costs will put you high on your brokers' holiday gift lists. Indeed, there are many large cheeseballs arriving at the Aronson+Partners trading room at year-end. But they could very easily be larger. The firm is very focused on effective trading. They have developed a particularly effective strategy, exploiting a wide range of traditional, electronic, and package trading tactics. Costs are closely monitored and used to inform future trading. They use "information from trading in trading and portfolio management" effectively to improve investment performance.

Could they do better? Most likely yes. More than 92% of the $15 million in shortfall costs was due to small trades, under 25% of average daily volume. There are just too many of them to work carefully without a higher level of automated assistance. If these small trades are as simple as everyone seems to think, this portion of costs can be reduced substantially.

SUMMARY

Equity markets and the tools used to participate in them have been transformed by technology. These changes affect all participants in the process. Investors have new ways to get their money's worth from equity management, and managers have new ways to transmit value. Trading is a critical link in the process.

There are repeated lessons in pictures of trading we see in all these trumpet-shaped charts, spanning ten years, and (thanks to all the outside correspondents) many managers and many countries. Large numbers of expensive, small "no-brainer" trades represent large potential savings, which goes directly to improved performance.

Intelligent "robotic" trading systems will have a hard time finding eight days volume of a stock, in a hurry. But they are ideally suited to sweeping multiple pools of liquidity, remaining cognizant of changing market conditions and the urgency of a particular trade. The large profits observed for many of these small trades underscore the potential for success in this area

A great deal of innovation is directed toward these goals. The relentless progress in technology means that traders really do have everything computationally, and they face an interesting set of choices in deciding where to put it.

ACKNOWLEDGEMENTS

I am grateful to many of my colleagues in the securities industry and academia for their helpful discussion and insights, both in the original and updated versions of this paper. Particular thanks to Ted Aronson, Greg Rogers, Mike Epstein, Ananth Madhavan, David Rothenberg, Richard Rosenblatt, John Dorian, Rob Arnott, Wayne Wagner, Erik Sirri, Andre Perold, Andy Lo, Dale Prouty, Yossi Beinart, John Mulheren, Ray Killian and Evan Schulman. I also thank the institutional brokerage community for all those cheeseballs.

REFERENCES

Bessembinder, H. and Kaufman, H., "A Comparison of Trade Execution Costs for NYSE and NASDAQ-Listed Stocks", Journal of Financial and Quantitative Analysis, vol. 32 no. 3, Sept. 1997, pp.287-310.

Bruce, Brian (editor), "Transaction Performance: The Changing Face of Trading", Institutional Investor Journals, Spring 2002

Chan, L.K., and Lakonishok, J., "Institutional Equity Trading Costs: NYSE vs NASDAQ", *Journal of Finance*, vol. 52, no. 2, June 1997, pp. 713-735

Collins, Bruce, and Frank Fabozzi. "A Methodology for Measuring Transaction Costs." *Financial Analysts Journal*, March/April 1991, pp. 27-36.

Domowitz, I., Glen J, and Madhavan, A., "Liquidity, Volatility, and Equity Trading Costs Across Countries and Over Time", InternationalFinance, Vol. 4, No. 1, Summer 2001, pages 221-256.

Grossman, Sanford J. *The Informational Role of Prices*. Cambridge, MA, MIT Press, 1989.

Keim, D. and Madhavan, A., "The Cost of Institutional Equity Trades", Wharton School, Rodney White Center Working Paper No. 1057, 1998

Leinweber, David. "Nerds on Wall Street." In *Blending Quantitative and Traditional Equity Analysis*, AIMR ICFA Continuing Education Series, 1994.

Leinweber, David, And Yossi Beinart. "A Little Artificial Intelligence Goes A Long Way on Wall Street." *Journal of Portfolio Management,* Winter 1995.

Leinweber, David, Peter Swank, and David Krider. "Evolutionary Ideas in International Style Management." In T. Daniel Coggin and Frank Fabozzi, eds., *Equity Style Management*. New Hope, PA: Frank Fabozzi Associates, 1995.

Leinweber, David, "Using Information from Trading in Trading and Portfolio Management", Journal of Investing, Summer 1995, pp. 40-50

Lo, Andrew, Craig MacKinlay, and Jerry Hausman. "An Ordered Probit Analysis of Transaction Stock Prices." *Journal of Financial Economics*, 31 (1992), pp. 319-379.

Loeb, Thomas. "Is There a Gift from Small Stock Investing?" *Financial Analysts Journal*, January/February 1991, pp. 39-44.

-----. "Trading Costs: The Critical Link Between Investment Information and Results." *Financial Analysts Journal*, May/June 1983, pp. 39-43.

Pérold, André F. "The Implementation Shortfall: Paper versus Realty." *Journal of Portfolio Management,* Spring 1988, pp. 4-9.

Pérold, André F., and Erik Sirri. 1997, "The Costs of International Equity Trading." Harvard Business School Working Paper 97-012

Treynor, Jack L. "What Does It Take to Win the Trading Game?" *Financial Analysts Journal*, January/February 1981, pp. 55-60.

Wagner, Wayne. "Transactions Costs: Measurement and Control." *Financial Analysts Journal*, September/October 1990.