

Market Impact

Market Impact is a very important component of investment returns for many funds. In the September edition of this newsletter, I showed that Claymore S&P/TSX CDN Preferred Share ETF (CPD) had a tracking error of 94bp versus its index in the twelve months to August 31, 2010; almost double its Management Expense Ratio of 48bp.¹ If we may assume that a negligible amount of this tracking error was due to mismatches between the fund's holdings and its index (the S&P/TSX Preferred Share Index, "TXPR") – which is an entirely reasonable assumption – then we may ascribe the entire otherwise unexplained shortfall of 46bp to market impact.

An example of one of the disruptions of normal relationships between issues which may be confidently attributed to Claymore's influence in the preferred share market was their purchase of POW.PR.C² to reflect a change³ in TXPR. This purchase was briefly discussed in the February, 2010, edition of this newsletter, but one of the charts from that issue is reproduced here as Chart 1 – a brief glance at the relative end-of-day closing bid-side yields for the two issues that should be very closely related will illustrate the market impact of this purchase.

One reason transaction costs were so high for that particular trade is that CPD is mandated to replicate TXPR – the fund is not permitted to substitute assets in the desired portfolio no matter how close the match, and delay in achieving the desired match is also not permitted.⁴ However, even in cases in which the portfolio manager has discretion to do what is best for the client, market impact is often either ignored or minimized via *ad hoc* adjustments to a model portfolio.⁵

For example, in the September, 2010, edition of this newsletter, I cited a 2004 study of trading by ABP, a Dutch pension fund ranking among the five largest in the world with assets of EUR 156-billion, which found that⁶ *average market impact costs equal 20 basis points for buys and 26 basis points for sells; average execution costs (defined as the sum of commission and market impact) equal 27 basis points and 33 basis points, respectively.*

Another estimate of market impact costs has been made by Calyon Financial, using data for the S&P500 eMini contract⁷: *Using data for the first quarter of 2006, we first calculated the distributions of sweep to fill prices at the beginning of each minute of each trading day. For each of these snapshots, we calculate five sweep-to-buy prices (exhausting the number of contracts offered at each price) and five sweep-to-sell prices (again, exhausting the contracts bid at each price). We then calculate the absolute value of the difference between these sweep-to-fill prices and the true market price at each moment. We then proceeded to see how many contracts actually traded in the interval immediately following each snapshot (five seconds in the case of E-mini S&Ps) and calculated the volume-weighted average price at which these trades were done.*

Comparing the two distributions, we find that actual trade prices reveal more liquidity than is apparent in the limit order book. As shown in Exhibit 4, the effect of hidden liquidity was worth slightly more than a cent for small orders and just under two cents for fairly large trade sizes. For intermediate-sized trades, though, the presence of hidden liquidity was worth considerably more. For trades between 500 and 1,000 contracts, hidden liquidity was worth 4.4 cents, while for trades between 2,000 and 3,000 contracts, hidden liquidity was worth about 3.8 cents per contract.

Table 1: Hidden Liquidity Summary for E-mini S&P 500 Futures 1Q06

Order Size	Average Impact		
	Sweep to Fill	VWAP	Difference
0 – 500	0.129	0.117	0.011
500 – 1000	0.189	0.145	0.044
1000 – 2000	0.243	0.202	0.042
2000 – 3000	0.332	0.295	0.038
3000 – 5000	0.428	0.411	0.017

Source: Calyon Financial

It should therefore come as no surprise that the issue of market impact costs has been studied intensively by both practitioners and academics and the concept has been extended to Implementation Shortfall (see Chart 2 for one example).

¹ See <http://www.claymoreinvestments.ca/en/etf/fund/cpd> (accessed 2010-11-12)

² See <http://www.prefblog.com/?p=9495>

³ See <http://www.prefblog.com/?p=9318>

⁴ Claymore, *Prospectus: Claymore ETFs*, available on-line at http://www.claymoreinvestments.ca/libraries/literature_en/prospectus_crq_clu_cic_cjp_cdz_cyh_cpd_clo_cmw_cww_cbq.pdf (accessed 2010-11-12)

⁵ Nicolò Torre, *The Market Impact Model™ – First in a series: The Market Impact Problem*, BARRA Inc. Publication No. 165, ("Horizon", Winter 1998), available on-line at <http://www.barra.com/newsletter/nl165/MIMNL165.asp> (accessed 2010-11-12)

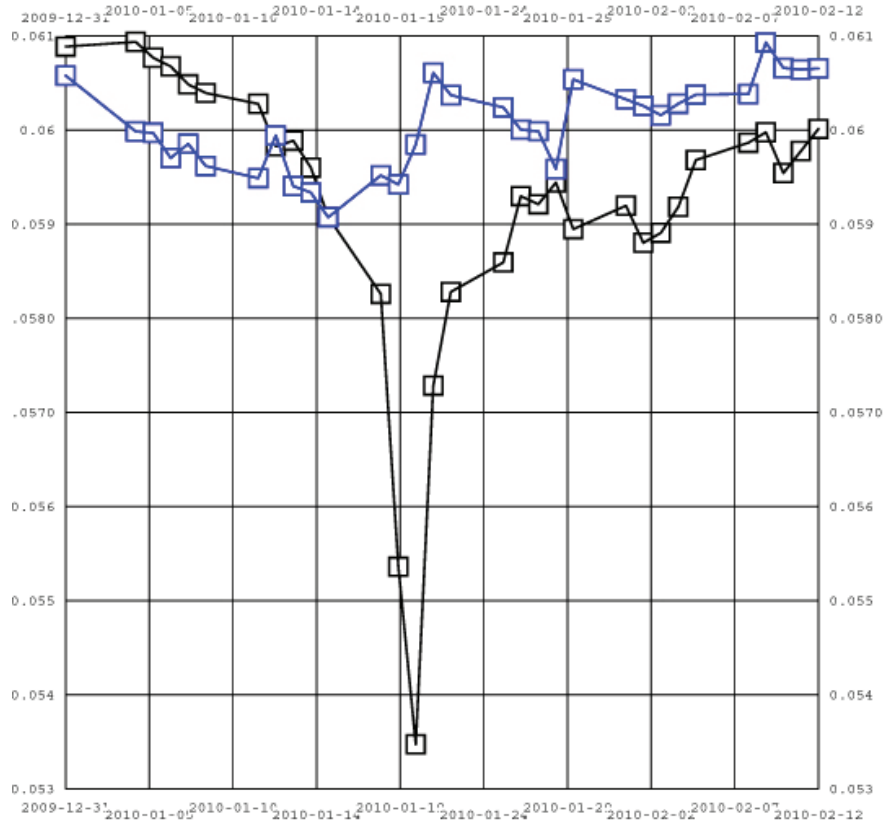
⁶ J. A. Bikker, L. Spierdijk and P.J. van der Sluis, *Market impact costs of institutional equity trades*, ISSN 0169-2690, June 2004, available on-line at <http://eprints.eemcs.utwente.nl/3545/01/1725.pdf> (accessed 2010-11-12)

⁷ Galen Burghardt, *Measuring Market Impact: Transaction Cost Analysis Comes to the Futures Market*, Futures Magazine, Nov–Dec 2006, available on-line at http://www.futuresindustry.org/downloads/fimag/2006/outlook07/outlook07_tca.pdf (accessed 2010-11-11)

1

POW.PR.C (Security A43503) Properties from 2009-12-31 to 2010-02-12

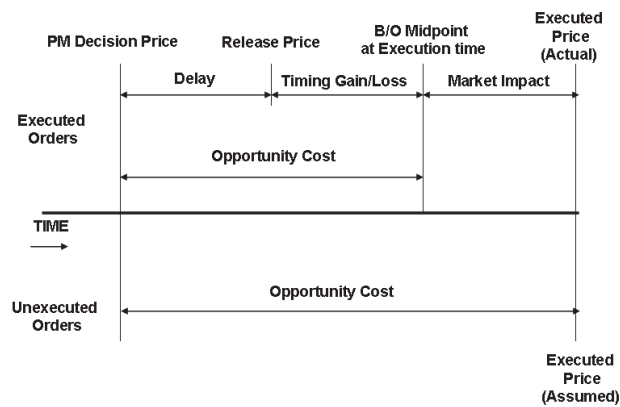
Yield-to-Worst (at Bid) Tax Identifier: 6
 POW.PR.B X-Axis: Date
 Y-Axis: Yield-to-Worst (at Bid)



Historical Market Data Source: TSE (c) 1993-2009 The Toronto Stock Exchange. All Rights Reserved

2

Types of Costs



In a 2003 presentation at Yale⁸, Zhiwu Chen made the point that there are enormous costs incurred due to delay between the time the Portfolio Manager makes the decision and the time the order is released by the portfolio manager's trading desk to the executing broker. The term "Implementation Shortfall" adds the costs due to organizational delay to the actual market impact costs; it should be noted that such delays might not be due simply to bureaucracy; a very large order might be executed over a period of days, with each piece of the order being released to the broker at the start of every trading day. Traditionally, the broker's efficiency in filling the order will be measured by comparing the average execution price to the Volume Weighted Average Price (VWAP) on a daily basis – in 2002, about 20% of US volume was traded to meet this benchmark price.⁹

However, it stands as an indictment of organizational chaos that he cites as potential benefits of a review of trading costs that:

- sometimes it might be considered justifiable to absorb costs in exchange for execution certainty, and
- Portfolio Managers should be more aware of the liquidity characteristics of their trades and potential costs.

The PMs are not managing portfolios of companies; they're managing portfolios of securities in those companies. For a PM to be unaware of the liquidity characteristics of their trades and potential costs is like a writer being unaware of grammar.¹⁰

The size of the manager also has a bearing on Implementation Shortfall:¹¹ *During 2Q10 average mega manager US equity Implementation Shortfall costs were +64.26 basis points as compared to +51.44 for all managers combined.*

For purposes of this essay, however, I will restrict myself to the discussion of market impact. If we assume that:

- Market impact costs are about 20bp¹² on each side of a trade
- USD 4,950-billion held in US domiciled equity mutual funds¹³
- A turnover rate of 64%¹⁴ (see Chart 3)

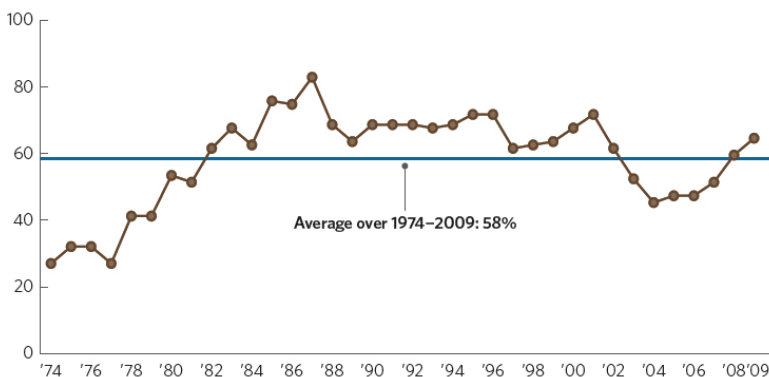
We may estimate that equity mutual funds in the US incur market impact costs of approximately USD 12.7-billion annually. Note that this calculation accounts for 'turnover' being comprised of one purchase and one sale, each side of which incurs the estimated 20bp market impact cost. This cost may be underestimated since it is taken from a pension fund for which cash flows will be both relatively small and predictable, unlike those of mutual funds. Mutual funds will also be constrained to take market action at times when market impact will be the highest.

3

FIGURE 2.6

Turnover Rate¹ Experienced by Equity Fund Investors²

Percent, 1974-2009

¹The turnover rate is an asset-weighted average.²Variable annuities held in mutual funds are excluded.

Sources: Investment Company Institute, Center for Research in Security Prices (CRSP), and Strategic Insight Simfund

⁸ Zhiwu Chen, Yale School of Management, *Transaction Costs and Trading*, October 1, 2003, on-line at <http://www.som.yale.edu/Faculty/zc25/HedgeFund/Trading-costs-lecture.ppt#257.1>, Transactions Costs and Trading

⁹ R. Scott Morris, *VWAP and Optimal Execution*, Goldman Sachs 2002-9-19, available on-line at http://premiacap.com/QWAFEFW/morris_20020919.pdf (accessed 2010-11-12)

¹⁰ Although that, too, seems fashionable these days.

¹¹ Elkins McSherry LLC, *Newsletter*, September 2010, available on-line at https://www.elkinsmcsherry.com/em/pdfs/Newsletters/Sept_2010_newsletter.pdf (accessed 2010-11-12)

¹² J. A. Bikker, L. Spierdijk and P.J. van der Sluis, *supra*

¹³ Investment Company Institute, *Investment Company Fact Book 50th edition*, 2010, ISBN 978-1-878731-48-3, available on-line at http://www.ici.org/pdf/2010_factbook.pdf (accessed 2010-11-12)

¹⁴ *ibid*

As I have previously remarked¹⁵, one man's market impact cost is another man's market making gain. USD 12.7-billion is a huge amount of money by any standards and, since it has been estimated using solely mutual funds, it may be assumed that these market impact costs can be captured by institutional market makers or other specialists who make a point of capturing this money.

In days gone by, it was necessary to be a specialist recognized by the exchange – or at least a knowledgeable floor trader – to have a chance to earn that money through the provision of liquidity, but those days are long gone by. Trading is done electronically and one may participate directly in the market without having a physical presence on the trading floor; additionally, order books have been made public, decreasing the advantage of the appointed specialist. We are now in the position where the competition to capture market-making profits has opened up to include entities other than formally recognized brokerages. And thus, we have the advent of High Frequency Traders (HFT).

Jonathan Brogaard of Northwestern University has estimated¹⁶ that High Frequency Traders realize gross trading profits of \$2.8-billion annually, about one-seventh as much as traditional market makers (who are more willing to hold inventory). Additional market-making profits will be earned by participants in block trades and opportunistic traders.

High Frequency Traders have become a favourite whipping boy for many market commentators and I believe that this is due to the fact that older, established firms and individuals, long accustomed to making easy profits from market-making, are now finding it harder to maintain their profitability. For instance, commentary from BMO Capital Markets in July, 2009, states¹⁷: *As more participants crowd out passive flow and lock markets to take advantage of the passive rebates it will become much more difficult for liquidity seeking investors to trade on the passive side of the market. This forces them to cross spreads and pay aggressive trading fees resulting in significantly higher trading costs for retail and institutional clients alike.*

This complaint simply reeks of an attitude of entitlement. I can think of no reason why a liquidity seeking investor should feel entitled to trade on the passive side of the market – it's always nice when you can, of course, but complaints that it is no longer possible should be given short-shrift.

In fact, it is my belief that HFT represent, at least to some extent, a triumph of merit over connections in the industry. The BMO-CM commentary noted above, for instance, complained that *We have had several discussions with Canadian buy side accounts who have noted the decreasing effectiveness of their pre-trade analytic tools. Typically these tools rely on volume and a number of other market metrics (e.g. spread, volatility) to predict the impact a given order will have on the market for that issue. As 'real' volume becomes less discernable these tools have greater difficulty determining this number. Portfolio managers, who have become increasingly reliant on these tools over the last several years, are becoming increasingly frustrated with their performance.*

These 'pre-trade analytic tools' do not sound very sophisticated to me. In the first place, it should be self-evident that these tools have been degrading in value with the passage of time and the increased competition in the marketplace – but the managers in question have not stepped up their game to account for either the new competition or the new analytical techniques the have become available. For instance, the order book (which shows not just the best bid and best offer, but additional levels of liquidity beyond these levels) contains useful information – which may be used not simply to estimate the cost of processing a market order (as shown in the Calyon Sweep to Fill calculations discussed above¹⁸) but also serves to improve predictions regarding the very short-term future average trade execution prices of the security in question.¹⁹

A Digression on Trading Strategy

In addition, I will mention that I am highly suspicious of any trading strategy that needs to include an allowance for any market impact at all; in most cases market impact, as defined by the difference between desired price and execution price, is simply an unnecessary cost that degrades performance.

The reason for this is that most portfolio activity in the market does not (or should not) result in any significant cash flow, since most trades are entered into, implicitly if not explicitly, as a swap between two issues. One does not (or should not) sell security A to generate some cash and then invest this magically appearing cash in security B: most of the time, a portfolio manager will wish to sell security A to buy security B.

Thus, it is the spread between the two execution prices that should concern the PM (and his trading team, if he has one), not the two execution prices considered separately. Further, by making a decision that the swap is in his clients' best interests, the PM is implicitly assigning a spread between the two issues to serve as a desirable range; and it is equally clear that there is an upper limit to the spread, above which the trade is no longer in his clients' best interests (or such that evaluation of the potential reward of trade execution is sufficiently unclear that it is outweighed by the known costs).

Further, it should be obvious that in the absence of new information, the spread between the two issues has improved since the last time the trade was evaluated; since if the spread has been constant (or moved against the trade) the trade would have been executed at the time of that prior evaluation. New information will certainly cause the desired spread to change, but, I submit, not very much, and not so often as prices change. If it was solely new information that caused price changes, these would alternate between long periods of minimal changes punctuated by sudden jumps as genuinely new and relevant information reached the market, which is not observed.

Thus, when trading in the absence of significant portfolio cash flows, a competent PM will specify a price spread (or, to be more precise, a price ratio) between the two issues traded and this specified spread will not be very different from the spread at the time that execution is first attempted.

¹⁵ September, 2010, edition of this newsletter

¹⁶ Jonathan Brogaard, "High Frequency Trading and its Impact on Market Quality", 5th Annual Conference on Empirical Legal Studies Paper, 2010-7-16, available on-line at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1641387

¹⁷ BMO Capital Markets Quantitative Execution Services, *The Impact of High Frequency Trading on the Canadian Market*, July 22, 2009, available on-line at http://qes.bmocm.com/papers/8_BMO_ImpactOfHFT.pdf (accessed 2010-11-13)

¹⁸ Galen Burghardt, *supra*

¹⁹ Charles Cao, Oliver Hansch and Xiaoxin Wang, *The Informational Content of an Open Limit Order Book*, available on-line at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=565324&http://papers.ssrn.com/sol3/papers.cfm?abstract_id=565324 (accessed 2010-11-13)

In this way, all considerations of market impact become moot. The PM is no longer saying “Buy 100,000 A and sell 100,000 B”; he is attempting to “Buy A and Sell B; equal numbers of shares on each side, as many as possible up to 100,000 on each side; at a maximum price spread of X.” If it turns out he can only get 5,000 shares of this trade done ... well, that’s too bad, but perhaps markets will be different tomorrow. Perhaps A will get cheaper; perhaps B will get more expensive; or perhaps C will drop suddenly in price due to a large sale and a different trade may be attempted.

But, disappointing as it may be to have only executed 5,000 shares of the swap, why would a competent PM wish to execute more? If the spread has moved beyond his desired spread, then by definition execution will not be in his clients’ best interests, as determined by him in good faith to the best of his abilities. Why would one want to execute such a trade?

I contend that this is a huge relative advantage that “real money” traders have over market-makers and HFT; an ability that is not sufficiently exploited. The competition is under great pressure to square their books by the end of every closing day; accumulating an inventory in the morning means that the afternoon’s trading will have a strong directional component, even if the market changes unfavourably. In contrast, a “real money” trader, having accomplished half the desired swap in the morning can simply walk away from it if the trade moves away in the afternoon – disappointed at not having completed the full desired amount, but happy that half got done. This advantage is lost – or implemented only on an ad hoc basis – with trading strategies that require the computation and forecasting of market impact.

A Further Digression: Spread-VWAP

Given that the economic purpose of a sale is, generally speaking, not to generate cash but to fund the purchase of something else, it should be clear that VWAP is a flawed measure of trading efficiency because it does not seek to measure economic utility.

For instance, let us consider two securities, A and B, which are priced equally. The portfolio manager has determined that selling A to buy B is economically desirable, provided that the trade is executed at the same price on each side of the trade. Further assume that the prices are very close to each other throughout the trading day, but that volume increases throughout the day for security A while decreasing for security B. Finally, assume that the price of both securities declines throughout the day.

It should be apparent that if the trade is executed as two separate legs, with the day’s VWAP being achieved for both sides of the trade, then the average price received for A will be less than the average price paid for B and hence that the trade as executed was not economically desirable. However, when the PM complains to his trading desk, he is met with a blank stare and reminded that execution was achieved with no deviation from VWAP; therefore the trading desk did a perfect job.

A performance metric that estimated the amount of trades that could have been performed during the day at the desired spread and gauging the trading desk’s efficiency by what proportion of this theoretical maximum it was able to capture would be much more economically relevant.

Even More Digression: Market Impact Costs Incurred by MAPF

As readers will know, I manage the Malachite Aggressive Preferred Fund.²⁰ The fund is small, but frankly, trading it is not much different from the trading I did with \$1.7-billion in government bonds while with Greydanus Boeckh & Associates.²¹ A universe of related securities is defined; a model of securities prices within this universe is developed; a trading engine is developed and calibrated to determine when a trade has sufficient chance of reward to make it worth the risk and costs; and, finally, when a trade for a live portfolio is identified, execution is attempted within the spread identified by the trading engine. If the trade can be executed – great! If the trade cannot be executed – walk! It’s not the trades you execute that create superior performance – it’s the trades you walk away from that matter.

All trades executed by MAPF in 2010 were isolated and the trading price compared with the closing bid price on the day of the trade in an attempt to quantify the market impact of a trade. For example, on October 22 the fund’s records show a purchase of 2,000 shares of GWO.PR.L at a price of 24.96, whereas the closing bid on that day was 24.90. Hence, the market impact for that trade was \$0.06 and the total cost of execution for that trade, prior to commission, was \$120.

For the year to November 5, there were 757 entries in the table²², with a total of 574,900 shares traded. It was with some surprise that I found that, even with same-day mark-to-market at the bid price, the fund made a profit on market impact of \$8,277.80 about 1.4 cents per share. Frankly, I had expected to see some market impact cost, but a profit is acceptable.


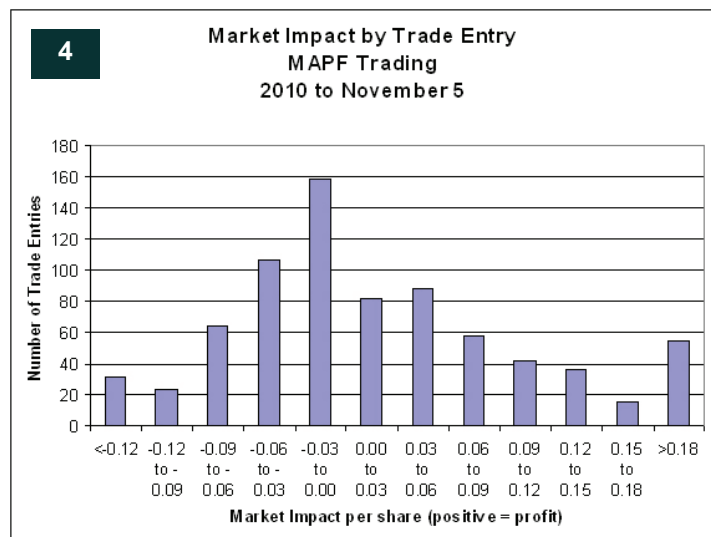
Note that the profit, while appreciated, is negligible – the fund’s outperformance versus its benchmark year-to-date is an indication that security selection was sufficiently accurate that on days following the trade date, the securities bought did better, in aggregate, than the securities sold.

Chart 4 shows the market impact by trade entry (equally weighted for frequency counting purposes, but ranging between 100 and 2800 shares) and are aggregated by trading day in Chart 5. Outlying values in the former chart are most often due to trades being performed at an extreme point in the intra-day fluctuations; since virtually all trades are entered into as part of a swap, the high values are more or less balanced by low values. Outlying values in the latter chart are due mainly to days on which the fund had cash-flows and therefore had purchases (sales) for which there was no complementary sale (purchase).

²⁰ See <http://www.himinvest.com/malachite/MAPFMain.php> for more information

²¹ Until its purchase and destruction by TD Asset Management

²² Trades are recorded as the fill is reported to me. A single order may give rise to many bookkeeping entries.



The Flash Crash: Overview

The SEC published their report on the Flash Crash²³ on September 30, summarizing events as follows:

On May 6, 2010, the prices of many U.S.-based equity products experienced an extraordinarily rapid decline and recovery. That afternoon, major equity indices in both the futures and securities markets, each already down over 4% from their prior-day close, suddenly plummeted a further 5-6% in a matter of minutes before rebounding almost as quickly.

Many of the almost 8,000 individual equity securities and exchange traded funds (“ETFs”) traded that day suffered similar price declines and reversals within a short period of time, falling 5%, 10% or even 15% before recovering most, if not all, of their losses. However, some equities experienced even more severe price moves, both up and down. Over 20,000 trades across more than 300 securities were executed at prices more than 60% away from their values just moments before. Moreover, many of these trades were executed at prices of a penny or less, or as high as \$100,000, before prices of those securities returned to their “pre-crash” levels.

By the end of the day, major futures and equities indices “recovered” to close at losses of about 3% from the prior day.

The report traces all the blame to “a large fundamental trader” that initiated an algorithm to sell 75,000 S&P-500 E-Mini futures contracts (with a notional value of about \$4.1-billion). The SEC claims that this algorithm was set “to target an execution rate set to 9% of the trading volume calculated over the previous minute, but without regard to price or time.” As will be seen in the following sections, this sentence appears to have been written by lawyers – it gives a hopelessly inaccurate description of the algorithm without being demonstrably false.

The initial selling pressure on this contract was initially absorbed by HFTs, some fundamental buyers, and arbitrageurs who opportunistically purchased contracts while hedging their position through sales of the ETF SPY and individual equities.

However, once the HFTs reached their position limits at about 2:41pm, they aggressively sold their positions, with a lot of trading being executed between them in an effort to dump the “hot potato” of a long E-Mini position. The SEC states that this increased volume caused the Sell Algorithm to increase the rate at which orders were entered into the market. In the four minutes from 2:41pm to 2:44pm (inclusive), the price of the E-Mini declined by approximately 3%.

This sudden drop and the absence of opportunistic buyers created further incentive to dump contracts; between 2:45:13 and 2:45:27, the price dropped another 1.7% to reach the day’s low of 1056. In the four-and-one-half minutes from 2:41pm through 2:45:27pm, prices of the E-Mini had fallen by more than 5% and prices of SPY had declined over 6%; trading in the E-Mini was halted for 5 seconds and the market commenced its recovery. At the time of the trading halt, the Sell Algorithm had sold about 35,000 E-Mini contracts, just under half the intended total, with a notional value of about \$1.9-billion.

There are, of course, many details to support this narrative, but the SEC report casts the events as a simple liquidity crisis caused solely by the Sell Algorithm as the driving factor. But there are problems.

²³ Findings Regarding the Market Events of May 6, 2010, Report of the Staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues, available on-line at <http://sec.gov/news/studies/2010/marketevents-report.pdf> (accessed 2010-11-13)


The Sell Algorithm was Better Behaved than the SEC States

The mutual fund company Waddell & Reed was identified very early as the originator of the 75,000 contract Sell Algorithm and, following the release of the SEC report was eager to clear its name.

Accordingly, it arranged with its executing broker, Barclays, to provide an independent company, Nanex, with data regarding the execution of the algorithm.²⁴ Nanex is a small company providing “ticker plant” technology (data feeds from the exchanges) to the industry, that has received considerable respect through its initial analysis of the Flash Crash,²⁵ documentation of data delays in the NYSEs dissemination of quote data,²⁶ publication²⁷ of “crop circles” (strange sequence of orders, repeated and cancelled very quickly), and continued investigation²⁸ of odd high-frequency quoting behaviour.

Nanex concludes²⁹ that *the algorithm was very well behaved; it was careful not to impact the market by selling at the bid, for example. And when prices moved down sharply, it would stop completely.* Nanex places more weight on the behaviour of the intermediaries who initially purchased these contracts: *The sale was so furious, it would often clear out the entire 10 levels of depth before the offer price could adjust downward.*

Note that the Nanex conclusion does not quite give the lie to the SEC’s description: the SEC’s use of the word “target” when describing the trading rate does not necessarily mean that the target volume was not adjusted or overridden within the algorithm; additionally, when the SEC states that the algorithm operated without regard to price, they are referring strictly to the lack of a limit price on the trade and ignoring the effect of price on the algorithm’s behaviour³⁰. The SEC seems determined to cast the algorithm in the worst possible light. This may be for political reasons: CFTC Commissioner Bart Chilton is openly hostile to what he characterizes as “robotic algo-trading”, while Chairman Gensler is eager to impose new obligations for brokers executing algorithms for their clients.³¹



²⁴ Graham Bowley, *Casting Doubt on Single-Trade Spark in Flash Crash*, New York Times Dealbook, 2010-10-8, available on-line at <http://dealbook.nytimes.com/2010/10/08/casting-doubt-on-single-trade-spark-in-flash-crash/> (accessed 2010-11-13)

²⁵ *Nanex Flash Crash Summary Report*, Nanex, available on-line at <http://www.nanex.net/FlashCrashFinal/FlashCrashSummary.html> (accessed 2010-11-13)

²⁶ See http://www.nanex.net/FlashCrashFinal/FlashCrashSummary_AlphaDelay.html and discussion at <http://www.prefblog.com/?p=12451>

²⁷ Nanex, *Analysis of the “Flash Crash”, Part 4, Quote Stuffing*, available on-line at http://www.nanex.net/20100506/FlashCrashAnalysis_Part4-1.html (accessed 2010-11-13)

²⁸ Nanex, *Daily High Frequency Quoting Stats*, available on-line at <http://www.nanex.net/HFTMonitor/HFTMonitor.html> (accessed 2010-11-13)

²⁹ Nanex, *May 6th 2010 Flash Crash Analysis, Final Conclusion*, 2010-10-14, available on-line at http://www.nanex.net/FlashCrashFinal/FlashCrashAnalysis_Theory.html (accessed 2010-11-14)

³⁰ CFTC Economist Andrei Kirilenk, quoted by Sarah N. Lynch, *Regulators Defend Report on ‘Flash Crash’*, Wall Street Journal 2010-10-13, on-line at <http://online.wsj.com/article/SB10001424052748703440004575548490476816252.html>

³¹ Sarah N. Lynch, *Regulator Calls for Curbs on Financial Futures*, Wall Street Journal 2010-11-1, available on-line at http://online.wsj.com/article/SB1000142405274870411104575588472618032264.html?mod=googlenews_wsj (accessed 2010-11-14)

The Role of Stop Loss Orders has not been Sufficiently Explained

The reason that the CME halted trading in the E-Mini contract at 2:45:28 was not directly due to concerns about contemporary trading; rather, the SEC report notes that *the CME Stop Logic Functionality was triggered to prevent the execution of the series of stop-loss losses [sic – I believe they mean “orders” – JH] that, if executed, would have resulted in a cascade in prices outside a predetermined “no bust” range.*

More importantly, however, retail stop-loss orders are dealt with sparingly in the report; primarily in the context of explaining why “internalizers” (companies that accept order flow from their own clients and other brokerages with a view to providing internal fills to many of these orders as an alternative to having them execute on a public exchange) sent orders to the exchanges at ridiculous prices, to be executed against “stub quotes”. The SEC notes the reason these orders were sent (essentially, the internalizers reached their position limits and had no capacity to purchase any more), but does not provide any specifics regarding the dollar value of these orders, or provide any discussion of their possible role as causes of the Flash Crash rather than effects.

The only estimate I can find of the volume of retail stop-loss orders was contained in a speech³² made by Mary Schapiro, Chairman of the SEC: *A staggering total of more than \$2 billion in individual investor stop loss orders is estimated to have been triggered during the half hour between 2:30 and 3pm on May 6.*

For the life of me, I do not understand why Ms. Schapiro chooses to cast those who entered these stop-loss orders as victims of the Flash Crash, rather than villains; I also do not understand why the effect of all this selling pressure was not dealt with thoroughly in the SEC report.

We know from the Waddell Reed/Barclays/Nanex data discussed above that the Sell Algorithm sold about \$2-billion notional worth of contracts prior to the crescendo of selling that took the market down to its day’s lows. And, from Ms. Schapiro’s speech, we can estimate that the value of sales generated by stop-loss orders was a further \$2-billion.

It is generally accepted³³ that the market impact cost of large execution increases on a per-share basis with the square root of the volume (at least as a first approximation), but it should be noted that this estimate is made during normal times in normal markets, not in times of great confusion when the market is experiencing a liquidity event. Thus, if you are prepared to estimate that the sale of 100,000 shares will cost you \$0.10 per share in market impact, your default assumption for the cost of selling 200,000 shares will be \$0.14 per share.

However, it should be noted that if this putative order for 200,000 shares is considered as two immediately sequential orders of 100,000 each, then the first 100,000 shares will take a charge of the originally estimated \$0.10 per share, while the second half of the trade will absorb costs of about \$0.18 per share, to bring average cost to the theoretically reasonable \$0.14.

At this point, I will refer to the data on prices from the earlier summary: *In the four minutes from 2:41pm to 2:44pm (inclusive), the price of the E-Mini declined by approximately 3%.*

This sudden drop and the absence of opportunistic buyers created further incentive to dump contracts; between 2:45:13 and 2:45:27, the price dropped another 1.7% to reach the day’s low of 1056.

These proportions are consistent with a hypothesis that:

- The initial selling by Waddell & Reed of about \$2-billion in notional value of the E-Mini contract took the market down 3%
- This price decline triggered retail stop-loss orders of a further \$2-billion (actual value)
- These stop-loss orders took the market down a further 1.7%. Further, unlike the E-Mini index contract, the effect of the stop loss orders was felt disproportionately by different securities, according to idiosyncratic volume and liquidity effects.

Note that this hypothesis differs from the Nanex conclusion³⁴ that *In summary, the buyers of the Waddell & Reed e-Mini contracts, transformed a passive, low impact event, into a series of large, intense bursts of market impacting events which overloaded the system.* However, Nanex is looking solely at the futures market data and does not investigate why the HFT players *slammed the market with 2,000 or more contracts as fast as they could.* It could very well be that HFT was reacting to the price effect of sell-orders on the equity market and was simply seeking to minimize losses by squaring their positions as quickly as possible. It is entirely reasonable to suppose that the HFT algorithms are informed by equity prices even if they do not actually execute arbitrage against them.

Allegations of Quote-Stuffing Have Not Been Investigated

Nanex documents³⁵ that the NYSE tape delays commenced very close to the time that the precipitous plunge in prices commenced at 2:42:44pm. It is possible that this was done to deliberate “quote-stuffing” for some unknown but presumably nefarious purpose; it is also possible that this was due to internalizers chasing the bid since upon receipt of a market order which they do not wish to fill, internalizers will place a limit order at the National Best Bid and Offer (NBBO); if it is not filled immediately, it will be cancelled and replaced with a new limit order at the new NBBO.

³² SEC, *Speech by SEC Chairman “Strengthening Our Equity Market Structure”*, Economic Club of New York, 2010-9-7, available on-line at <http://www.sec.gov/news/speech/2010/spch090710mls.htm> (accessed 2010-11-13)

³³ Esteban Moro, Javier Vicente, Luis G. Moyano, Austin Gerig, J. Doyne Farmer, Gabriella Vaglica, Fabrizio Lillo, and Rosario N. Mantegna, *Market impact and trading pro_le of large trading orders in stock markets*, and references therein, available on-line at http://arxiv.org/PS_cache/arxiv/pdf/0908/0908.0202v1.pdf.

³⁴ Nanex, *May 6th 2010 Flash Crash Analysis Final Conclusion*, 2010-10-14, available on-line at http://www.nanex.net/FlashCrashFinal/FlashCrashAnalysis_Theory.html (accessed 2010-11-13)

³⁵ Nanex, *Nanex Flash Crash Summary Report*, 2010-9-27 available on-line at <http://www.nanex.net/FlashCrashFinal/FlashCrashSummary.html> (accessed 2010-9-27)

I have not yet seen a satisfactory explanation for the huge influx of quotations that caused delays in the tape, however, I take heart from the SEC report's note that: *Accordingly, another area of focus going forward should be on the integrity and reliability of market centers' data processes, especially those that involve the publication of trades and quotes to the consolidated market data feeds. In addition, we will be working with the market centers in exploring their members' trading practices to identify any unintentional or potentially abusive or manipulative conduct that may cause system delays that inhibit the ability of market participants to engage in a fair and orderly process of price discovery.*

Stop-Loss Orders in General

A very puzzling feature of Ms. Schapiro's speech of September 7 is her apparent concern for the retail investors who placed stop-loss orders.

She emphasizes in her introduction that the social purpose of equity markets is to allow companies to raise capital from investors who know they may later wish to sell at a fair and efficient price and notes the importance of equity market trading as an indicator to the company of the levels at which it can raise funds in a secondary offering. It is important to all of us that individual companies be reasonably valued according to their fundamentals, so that capital can be deployed efficiently.

But a stop loss order has nothing to do with fundamental analysis. After all, a stop-loss order is a market order that gets triggered at a price worse than one that was effective when the order was placed. If you are willing to sell something at \$45, why are you not willing to sell at \$50?

When challenged on this point, one investor responded³⁶: *Fear & Greed*

If a security you own is in an uptrend, greed makes you want to keep it to see where it might go. However if it starts to go down, Fear makes you want to keep some of the profit. The answer is a stop loss order.

You don't sell at 50 because you are afraid it will go up and you will miss a move. Who knows perhaps your perception of the future value of the security is \$100.

You set 45 as a stop for various reasons. 45 is the maximum pain you can take, 45 is the break-even point, 45 is your original profit target, or you feel if the price get below 45 it will drop further allowing you to buy back at a lower price.

This has nothing to do with fundamental analysis – this is simply a view of financial markets as a momentum-driven casino. And yet the SEC chairman is sympathizing with the users of this type of order as victims in the same speech as she insists on the role of equity markets in fairly valuing corporations.

There is some speculation³⁷ that the SEC will introduce collars to all market orders; alternatively, this may be applied only to stop-loss orders, so that stop-loss market orders are banned and replaced exclusively with stop-loss limit orders. Similar in principal is the NASDAQ proposal³⁸ *to amend BX Rule 4751 to include Order Collar functionality that cancels any portion of any Unpriced orders (also known as market orders) submitted to the Exchange that would execute at a price that is more than \$0.25 or 5 percent worse than the national best bid and offer at the time the order initially reaches the Exchange, whichever is greater.*

However, there are many entrenched interest who are rather fond of stop-loss orders – Canadian discount brokerages include stop-loss functionality in their standard order screens and at least one promotes its use after every normal order is entered. Market Orders are very valuable to brokers – they can be sold to internalizers if the brokerage does not make markets itself – and represent an important source of revenue beyond their commission value. And finally, there is a huge pool of charlatans eager to convince the unwary that Stop-Loss orders are the key to huge risk-free profits. The strength of such faith should not be underestimated!

Conclusion

The Flash Crash appears to be simply a matter of market impact writ large. Regardless of whether the specific trigger for the plunge was the Waddell & Reed order to sell 75,000 E-Mini Contracts, an undisciplined squaring of books by the HFT intermediaries, or an avalanche of stop-loss orders after an initial decline, the bare fact appears to be that selling pressure momentarily overwhelmed the available liquidity provided by market makers, resulting in huge market-impact costs for those who sold – and, of course, equally huge market-making profits for those who were in a position to react quickly and effectively to the plunge.

Market Impact is a huge determinant of returns, but it should always be kept in mind that one man's market impact costs is another man's market making gain. Investors who are aware of these costs and take steps to minimize them will, generally speaking, outperform those who ignore such costs, or treat them simply as an adjunct to the investment management process rather than as an integral part of each trading decision.

³⁶ See <http://www.financialwebring.org/forum/viewtopic.php?f=33&t=1155&start=152> (accessed 2010-11-13)

³⁷ Sal Arnuk and Joseph Saluzzi, *What the SEC Staff Will Likely Recommend in Reaction to the May 6th Flash Crash*, Themis Trading, 2010-9-24, available on-line at http://www.themistrading.com/article_files/0000/0601/092410_WhitePaper.pdf (accessed 2010-11-13)

³⁸ SEC, *Release No. 34-63052; File No. SR-BX-2010-067*, 2010-10-6, available on-line at <http://www.sec.gov/rules/sro/bx/2010/34-63052.pdf> (accessed 2010-11-13)