Adverse Selection vs. Opportunistic Savings in Dark Aggregators

SERHAN ALTUNATA, DMITRY RAKHLIN AND HENRI WAELBROECK

SERHAN ALTUNATA is Research Analyst at Pipeline Financial Group, Inc.
Serhan.Altunata@pipelinefinancial.com

DMITRY RAKHLIN is Senior Quantitative Trader at AllianceBernstein.
Dmitry.Rakhlin@alliancebernstein.com

HENRI WAELBROECK is Vice President and Director of Research at Pipeline Financial Group, Inc.
Henri@pipelinefinancial.com

In the last few years we have seen an explosion in the number of small-order dark pools, leading to much talk about fragmentation. In fact this concern is misplaced: liquidity pools are connected to each other through direct links, by automated market makers that exploit arbitrage opportunities between displayed and dark markets, and through dark aggregators. Dark aggregators produce fairly consistent market participation rates of 15-25% of the total market liquidity, not unlike the rates most frequently used by the most popular trading algorithms. The average fill size of dark aggregators is also similar to that of conventional algorithms; so if we categorize trading methods broadly with block trading strategies on one hand and small-order strategies on the other, dark aggregators are best viewed as one of several algorithm styles that tap the liquidity of small-order markets. As such, when should traders expect them to perform better (or worse) than conventional algorithms? The intent of a dark aggregator is to exploit opportunities for "impact-free" liquidity away from the displayed-order markets; to understand how it performs we are led to confront the long-standing problem of how to weigh the potential gains of an opportunistic trading strategy against the cost of adverse selection.

In algorithmic trading, adverse selection refers to execution at an unfavorable price point shortly before a substantial price improvement opportunity, or conversely, a low participation rate prior to an adverse price move. Typically, one observes an increase in execution rate accompanied by a rise in the incremental trading volume in the market. After this brief activity, a trend develops, causing either remorse for having filled a trade too easily or frustration with the lack of execution as the stock moves away.

Being aware of adverse selection is especially important in high volatility markets. Mittal and Wong [2008] have shown that a cash-balanced version of an opportunistic algorithm strategy can result in significantly lower shortfalls than the underlying single-stock strategy itself – how is this possible? The effect of cash balancing is to reduce adverse selection by controlling the participation rate, so the shortfall reduction from cash balancing points to opportunities to improve single-stock trade performance by addressing adverse selection. Cash balancing in this sense is a first step towards full-fledged predictive control of the participation rate by managing the execution across algorithm styles [Stephens and Waelbroeck 2009].

If adverse selection plays such a key role in execution performance, are dark aggregators more exposed to adverse selection than other strategies? And if so, when is it profitable to use them?

To manage the tradeoff between shortfall and execution risk, the trader needs to know how much adverse selection costs in relation to other aspects of execution, such as market impact, spread or commission costs.

Historically, adverse selection has been defined only at the level of a single order, where the bid ask spread is a direct measure of adverse information risk as explained in [Glosten and Milgrom, 1985]. Since the idea of aggregate adverse selection for an algorithmic execution is clear only in concept, one can make a case for several methods to calculate it. The most extensive study on the subject has been recently published by Jeria and Sofianos [Jeria and Sofianos, 2009]. In this study, the authors argue that some adverse selection is natural and inevitable with passive strategies. They point out that execution costs are not only given by transaction costs, but there are implicit transactions costs associated with the lost short term alpha on the non-filled shares of an order[1]. The authors then quantify the adverse selection cost as a weighted average of shortfalls on filled and non-filled shares. Based on this measure, they conclude that a passive order in a dark aggregator is slightly less costly than a 20% participation strategy. While their approach improves over Mittal and Wong by explicitly considering clean-up costs, their conclusions rely on a definition of short-term alpha that is intertwined with the market impact of the dark strategy. More importantly, the method does not enable a separate measure of opportunistic savings and adverse selection costs. The ability to measure each of these two terms separately enables one to evaluate methods that tackle adverse selection without losing the potential for opportunistic savings.

Our goal ultimately is to help a trader decide when to use a dark aggregator as opposed to another type of opportunistic strategy, or when to use instead a low variance strategy which would be less exposed to the risk of adverse selection. This requires measuring adverse selection costs relative to opportunistic savings, to help the trader decide when the former outweighs the latter and choose strategies accordingly.

We take the view that short-term market dynamics is driven not by random order flow but by a competition between strategies that represent a deliberate, profit-seeking motive. The cost of adverse selection is the mirror image of the profitability of a strategy that takes the contra side to an adversely selected trade, and the opportunistic savings are someone else's adverse selection. If an opportunistic trading strategy is successful, it should achieve better-than-average prices. Vice-versa, an opportunistic strategy will be unsuccessful when it falls prey to traders with short-term alpha that exploit the implicit optionality in their order flow. In this case the cost of adverse selection would exceed opportunistic savings. The net average result of opportunistic savings minus adverse selection costs for an opportunistic trading strategy therefore measures its intelligence relative to the rest of the market.

A strong argument can be made that adverse selection in dark aggregators occurs when trading against informed or arbitrage traders who believe that the stock is mispriced. These traders sweep through dark pools to find resident midpoint liquidity before turning to the displayed markets. Therefore, it is quite plausible that increased execution rates in dark pools are lead indicators of short term market movements that affect the intraday P&L. This points to the potential for substantial benefits to be derived from methods that throttle the participation in dark pools and leverage the insights gained from the participation anomaly to guide the trading in lit venues.

In the first part of our study, we report on experiments carried out by Pipeline using its own capital to measure adverse selection by executing long-short sector baskets, holding the positions long enough for temporary impact to dissipate, then winding them down[2]. Leveraging the insight gained from these experiments, we define adverse selection in terms of execution performance relative to an imaginary 10% volume-tracking strategy, a benchmark known as the Participation Weighted Price (PWP) and first introduced by Abel Noser [see for example Jenkins 2001]. This benchmark strategy is evaluated using

---

[1] Jeria and Sofianos refer to these costs as clean-up costs. Here, we instead use the term opportunity costs.

[2] This is one major aspect of our methodology that differs from the methodology of Jeria and Sofianos; those authors use client data to measure short term alpha, which in turn is used to quantify adverse selection. However, the short-term alpha may be contaminated by potential other trades from the same client executing elsewhere and causing additional impact. By using well correlated and market neutral baskets as well as controlled sample of trades executed by a single institution, we remove such contamination from our data.

market data from the period where the actual trade was executed, so it incorporates the market impact of the actual trade but not the cost variance attributable to timing choices. Where the realized price is better than the PWP benchmark, the trader can claim that the opportunistic strategy chose its price points intelligently; in trades where the difference was unfavorable we conclude that the trade was subject to adverse selection. Taking a significant sample of random trades, we can calculate the average adverse selection cost and the average opportunistic savings of any trading method in absence of selection bias.

Of course institutional trades are not random; to verify that our conclusions apply also for real-life trades one of us (DR) reproduced the analysis at AllianceBernstein applying the same methodology to a large dataset of trades executed in a major agency dark aggregator from January to May 2009.

Our goal is to help traders enhance alpha capture by using opportunistic strategies tactically at times when they can be expected to boost an execution strategy, while avoiding them by switching to more risk-averse trading styles in situations where they are vulnerable to adverse selection. To evaluate the potential for this level of execution control, we consider a control sample of similar basket trades in Pipeline's Algorithm Switching Engine. Our main results can be listed as follows:

Main Results

- Opportunistic savings in dark aggregators are almost entirely lost to adverse selection.
- Approximately half of the implementation shortfall on completely-filled market orders can be attributed to intra-trade adverse selection.
- Simple measures employed by AllianceBernstein to control the participation rate in dark aggregators were able to significantly reduce adverse selection, lowering implementation shortfall relative to the "control" strategy.
- Predictive switching can eliminate 70% of the adverse selection costs with only a small reduction in opportunistic savings, resulting in a 40% lower implementation shortfall relative to continuous use of a dark aggregator.
- Our methodology is robust to selection bias and other forms of data contamination in real-world institutional trades.

Definitions of Adverse Selection and Opportunistic Savings

We will compare the implementation shortfall with respect to a 10% PWP benchmark, which is the interval VWAP for a period of time starting at the same point as the actual trade and carrying forward until 10 times the traded quantity has printed to the tape. This benchmark price reflects most of the market impact from the actual trade, because it is measured in a period where the trade was being executed - but it excludes any contribution to the shortfall that may result from variations in the trade schedules.

We define $IS_{10\%} = \ln(\frac{PWP_{10\%}}{S_{1i}})$ as the hypothetical shortfall of a 10% participation strategy. The tracking error relative to this benchmark represents the difference between the actual trade and the participation strategy,

$$IS_1 = IS_{10\%} + TE \qquad [1]$$

If the tracking error is negative, this implies that the algorithm is outperforming the benchmark trade and making smart decisions to reduce the implementation shortfall. In this case, the negative tracking error can be called *opportunistic savings*. In the opposite case, a positive value of the tracking error implies that the algorithm is underperforming by executing at worse than fair prices. The flip side to opportunistic savings is *adverse selection*.

Taking an ensemble of trades, some produce opportunistic savings and zero adverse selection, while others are adversely selected in which case the opportunistic savings are null. We can then define adverse selection as:

$$\begin{aligned} AS &= \frac{1}{N}\sum_{k=1}^{N}\Theta(IS_1^k - IS_{10\%})(IS_1^k - IS_{10\%}) \\ &= \frac{1}{N}\sum_{k=1}^{N}\Theta(\ln\frac{P^k}{P_{Part}^k})\ln\frac{P^k}{P_{Part}^k}, \end{aligned} \qquad [2]$$

where $\Theta(x)$ is the step function which is zero for negative x and one for positive x. The values P^k refer the interval VWAPs obtained by the test algorithm and the values P_{Part}^k are the VWAPs from the benchmark participation algorithm.

In the same spirit, we can define the opportunistic savings as:

$$\begin{aligned} OS &= \frac{1}{N}\sum_{k=1}^{N}\Theta(IS_{10\%} - IS_1^k)(IS_{10\%} - IS_1^k) \\ &= \frac{1}{N}\sum_{k=1}^{N}\Theta(\ln\frac{P_{Part}^k}{P^k})\ln\frac{P_{Part}^k}{P^k} \end{aligned} \qquad [3]$$

The variation of AS and OS across execution strategies can shed light on the relationship between opportunistic savings and adverse selection. In particular, we can address the key question of whether algorithm designs and dark aggregators provide a net benefit to the trader given the tradeoff between adverse selection and opportunistic savings.

The average tracking error for a set of trades can be expressed in terms of the two components in Eqs.4 and 5 as:

$$TE = AS - OS. \qquad [4]$$

Adverse Selection and Intraday P&L in Unbiased Random Trades

Pipeline engages in random closed-cycle trades daily as part of its ongoing program to keep the Algorithm Switching Engine in tune with market conditions and incorporate new algorithms in the platform. The effects of adverse execution (exemplified in exhibit 1) are reflected directly in the daily profit and loss of these trades. This has heightened awareness of the problem and helped shape the design of the Engine. For institutional trades, adverse selection is hidden in delay and opportunity costs that are not so easily quantified: the desired execution size and limit prices are related to investment objectives specific to each trade; these objectives themselves change day to day as portfolio managers adjust their intentions in light of new information. For example, how is one to evaluate the opportunity cost for not completing a 200,000-share segment of a million-share buy trade that is expected to execute in the course of a week, when the stock moves higher the next day and retraces below arrival price by the end of the week? The answer becomes intertwined with the trader's view on likely price scenarios and the portfolio manager's instructions regarding the alpha and urgency in the trade. However, even though it is hidden, the same adverse selection cost exists in institutional trades as is observed in the daily profit and loss when testing algorithms. It is important to quantify adverse selection on a firm mathematical basis and to develop trading applications accordingly in an effort to maximize the efficiency of the execution process at whatever speed the trader chooses to work the order.

This study is based on test trades that were carried out from 11/19/2008 and 12/17/2008 with baskets of symbols specific to each sector and market cap grouping. Simultaneous long and short positions were taken across an even number of symbols to filter out any market effects. The basket composition was further optimized to minimize the variances of the basket returns based on the correlation matrix, subject to expected trade durations of at least 15-30mins. This makes our sample truly random as no alpha seeking strategies were employed. The dark aggregator and the switching engine were tested on alternating days with 181 completed trading segments, corresponding to a deployed capital of 11.2million USD. Starting 15 minutes after completion of the last trade, all positions are wound down using the same trading method.

Exhibit 1

Adverse selection occurs when arbitrageurs sweep dark pools then displayed markets to maximize liquidity capture ahead of a price move.

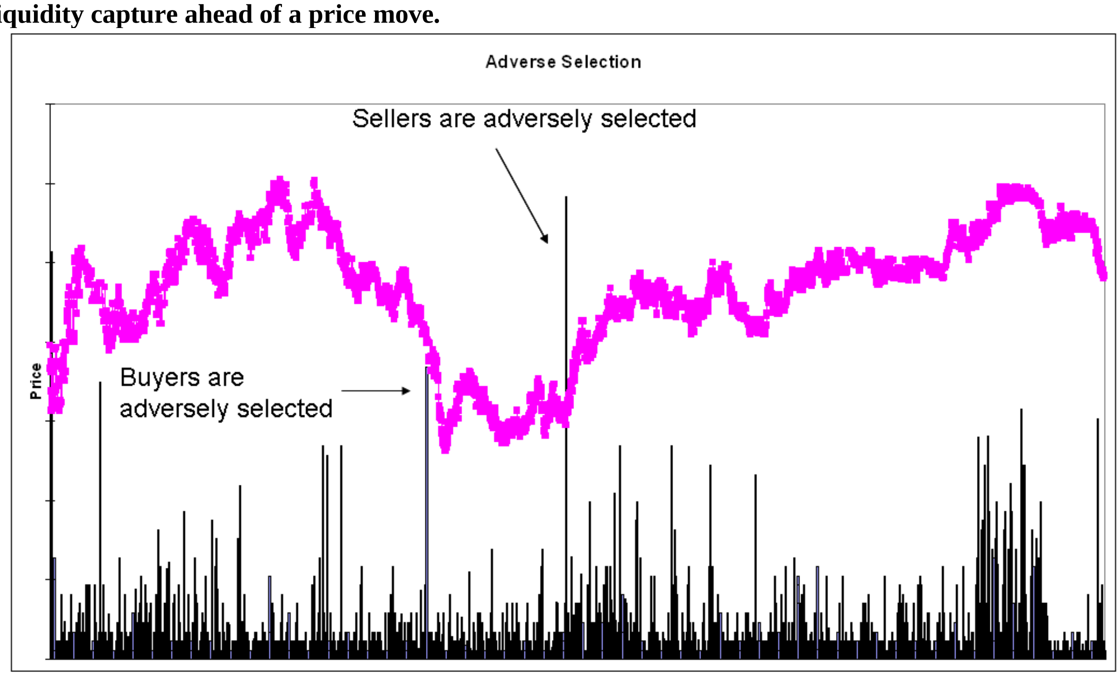

The P&L per share of a complete cycle starting with a buy trade is

$$PL = \ln(S_{2i} / S_{1i}) - IS_1 - IS_2\,(bps). \quad [5]$$

where S_{1i} is the price at the beginning of the entry segment, S_{2i} is the price at the beginning of the exit segment, and IS_1, IS_2 are the realized implementation shortfalls incurred on entry and exit respectively. The unrealized P&L after reversion following the entry trade is given by

$$U = \ln(S_{2i} / S_{1i}) - IS_1. \quad [6]$$

By definition,

$$\ln(S_{2i} / S_{1i}) = PI_1 = \alpha_{1,2} + E(PI_1)\,, \quad [7]$$

where $E(PI_1)$ is the expected permanent impact caused by position entry and $\alpha_{1,2}$ is the price return from start of entry segment to the beginning of the exit segment that is not attributable to the trade. In our case α has zero mean since we are executing random trades.

Using a pre-trade model to compute the expected shortfall of a 10% participation strategy, we can write $IS_{10\%} = E(IS_{10\%}) + \overline{\alpha}_{10\%}$ where $\overline{\alpha}_{10\%}$ is the average alpha for the 10% participation period, The unrealized P&L becomes

$$U = (\alpha_{1,2} - \overline{\alpha}_{10\%}) + (E(PI_1) - E(IS_{10\%})) - TE \quad [8]$$

The first term measures the incremental alpha from the 10% benchmark interval to the beginning of the exit trade; this is zero on average for random trades but if we measure unrealized P&L at the close of the 10% benchmark interval this term drops out on individual trades. The second term measures the difference between expected permanent impact and the expected implementation shortfall of the benchmark trade, which is zero on average if the market is efficient to information arbitrage [Farmer, Gerig, Lillo, Waelbroeck 2009]. Therefore we expect that the average unrealized P&L at the start of the exit trade is directly measured by the tracking error $U = -TE$. This highlights the importance of adverse selection in entry trades.

For example if a stock costs $20, an institution buys a large position with an expected impact of 100 bps. In average market conditions the price moves up to a peak of $20.30, with an average fill price of $20.20. After reversion the price falls back to $20.20, and the institution breaks even if the position is marked to market. In an over-supplied market, the trader using a passive strategy may find that the trade executes more quickly than expected at or very close to the arrival price of $20; after the trade is done the stock falls to $19.90 and the position is in the red: in this case, what was apparently a well executed trade is in fact losing value due to adverse selection.

Adverse Selection in Dark Aggregators

The results from test trades using one of the popular dark aggregators and the algorithm switching engine are summarized in Exhibit 2. These graphs display the distribution of adverse selection values for the dark aggregator and engine trading methods. The dark aggregator distribution is skewed with outlier values lying as far as adverse selection costs of 160bps. On the other hand the switching engine distribution is compact around zero with few outliers. This implies that adverse selection could be a costly component to trading in a dark aggregator. At the same time, the switching engine data shows that predictive switching and limit price management can reduce these adverse selection costs.

Exhibit 2
Adverse selection in a dark aggregator, compared to predictive switching

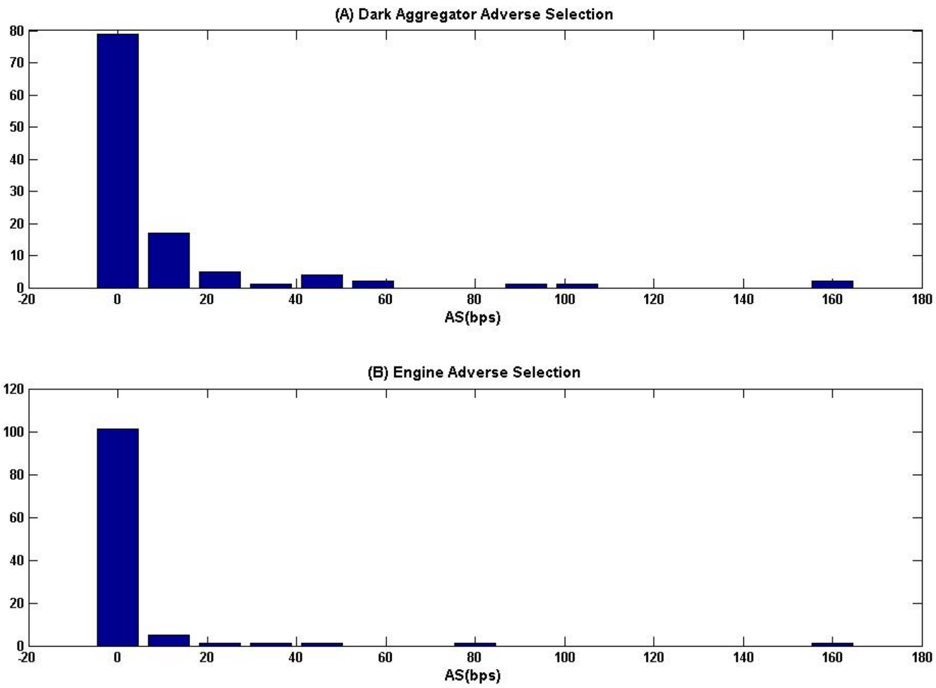

Opportunistic Savings in Dark Aggregators

We can also compare the opportunistic savings from the two methods of execution. Exhibit 3 shows the distribution of opportunistic savings.   Similarly, the dark aggregator shows opportunistic savings with more outliers than the Engine.

Exhibit 3
Opportunistic savings in a dark aggregator compared to predictive switching

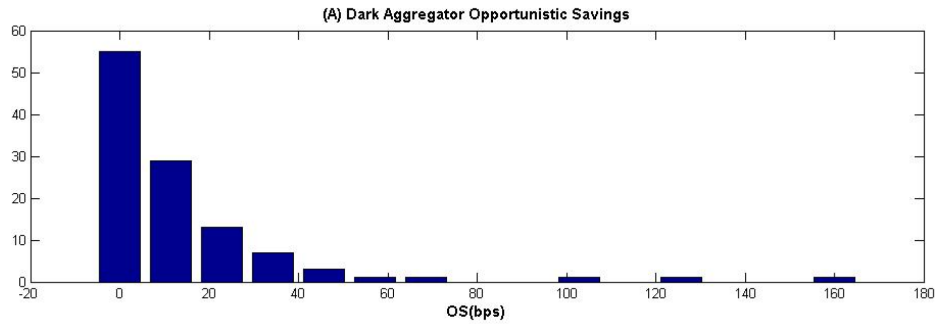

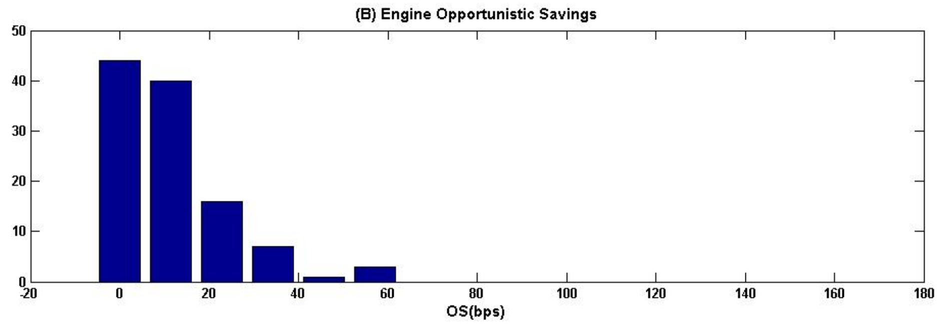

The results from our tests are summarized in Exhibit 4. The difference in adverse execution costs between the dark aggregator and switching engine are interesting because the switching engine itself uses dark aggregators, albeit tactically at specific points in time and with brief exposure times and tactical limit prices. The switching engine results therefore can be viewed as a proxy for the performance gains that a trader might be able to achieve also by managing tactical exposure to dark aggregators based on her intuition of market dynamics. The dark aggregator yields an average adverse selection cost of 11.2bps as compared to an adverse selection cost of 3.6bps in the engine. This signifies close to a 70% reduction in the adverse selection costs.

The opportunistic savings in the dark aggregator are 13.3bps, whereas the engine provides opportunistic savings of 11.8bps. The adverse selection in the dark aggregator is approximately 50% of the observed implementation shortfall.

What is the effect on the bottom line of a potential reduction in adverse selection? Based on the comparison given above, the 70% reduction in adverse selection in the switching engine reduces implementation shortfall over the dark aggregator by 40%. This stems from the fact that the switching engine uses a prediction engine for switching, which helps avoid the risks of using opportunistic trading styles at times where exposure to adverse information is higher. For example, if the prediction engine working a buy order foresees a developing up-trend it will refrain from using dark aggregators or passive strategies that would be likely to underperform, and choose instead low rate variance strategies.

Likewise, high touch traders can add value by using low variance algorithms in trending market conditions and seeking price improvement with passive strategies only when they anticipate mean-reverting markets. Ideally, algorithm vendors should be able to assist the trader by reporting on the mean and variance of the participation rates of their algorithms – for high variance (opportunistic) algorithms, the degree to which they are exposed to adverse selection can be measured as described in this paper. If upon carrying out this analysis for an opportunistic algorithm the trader finds that adverse selection costs exceed opportunistic savings, the algorithm could be better described as someone else’s opportunity.

Exhibit 4
Comparison of Adverse Selection Costs and Opportunistic Savings

Method	*Total Capital Traded*	*Adverse Selection*	*Opportunistic Savings*	*Implementation Shortfall*	*-TE*[3]
Dark Aggregator	5.6million USD	11.2bps	13.3bps	21bps	2.1bps
Algorithm Switching Engine	5.6million USD	3.6bps	11.81bps	12.5bps	8.2bps

Application to Institutional Trades

Having a sample of random, equally-weighted, correlated trades helps eliminate various selection biases inevitably present in real-life trading. In this section we show that the same methodology can be successfully applied to the analysis of institutional trades and that it is robust to demonstrate the same results with respect to a different dark aggregator.

We have analyzed two large samples of broker placements. The first sample covers the period of January-February, 2009 and consists of 1225 trades executed in one of sell-side dark aggregators and 2436 "control" trades. To form a control set we took non-aggregator placements in symbols that also had an aggregator placement on the same day. While most of the placements in the control set also had access to various dark pools, the aggregator was the only classical example of passive dark strategy:

- It could only execute in the dark. There were no dynamical rules to alter the pace of execution in the dark besides setting static limit prices and minimum crossing sizes on child orders.
- It had access to more dark destinations, including most broker pools.

The table below gives more stats on the trades in the sample; all columns except # orders are dollar-weighed averages. Executions in the same symbol do not overlap in time.

Exhibit 5
Comparison of dark aggregator and control sample of broker algorithms shows similar performance

	Orders	Traded $, bn.	Executed Size	Duration minutes	%Daily Volume	Participation Rate	%Filled	E(Cost) bps	IS bps	IVWAP bps	Close bps	Total Cost bps
AGGREGATOR	1225	1.0	155,000	18.4	1.2%	22%	91%	23	16.0	4.7	18	16.2
CONTROL	2436	1.9	141,000	21.6	1.2%	25%	95%	19	14.1	1.7	14	13.4

Trades in the aggregator and the control sets were very similar in size and expected cost. Only large- and mid-cap trades were included. Participation rates were also similar. Judging by realized implementation shortfall we cannot conclude if the aggregator performed any better or worse than the control set. Alpha to close was almost identical, and the total cost (Jeria and Sofianos, 2009 methodology) was within a basis point from IS. The aggregator's performance versus interval VWAP was better by 3 bps, but it is unclear why. The reason for both aggregator and control strategies to beat the expected cost can be explained by the fact that cost model input parameters (volatility, spread, etc.) are usually lagged, and in this case are partly taken from highly volatile 4Q 2008. In short, the only visible benefit to use the dark aggregator was flow diversification.

In February we worked with the broker to improve the aggregator's performance. The new strategy deployed to our desk in March had the ability to

[3] Negative values of the tracking error correspond to savings while positive values imply costs.

- Maintain execution within minimum and maximum participation rate range.
- Dynamically set limit prices on child orders based on the market conditions.
- Go the public venues if no liquidity can be found in the dark.
- Temporarily withdraw from dark destinations if gaming was suspected.

Below are the performance stats for the new strategy versus the control set for the period of March, 1 – May, 15, 2009:

Exhibit 6

Modified aggregator shows improved performance over the control sample

	Orders	Traded $, bn.	Executed Size	Duration minutes	%Daily Volume	Participation Rate	%Filled	E(Cost) bps	IS bps	IVWAP bps	Close bps	Total Cost bps
AGGREGATOR	1269	1.4	121,000	21.2	1.1%	21%	91%	18	7.2	4.6	22	6.7
CONTROL	2168	2.2	149,000	17.2	1.3%	31%	93%	20	11.5	2.7	22	11.8

The new control set shows only a moderate 15-20% decrease in IS compared to the January – February period, comparable with the decrease in expected costs due to lower volatility and tighter bid-ask spreads. The aggregator performance improved a lot more drastically. The chart below shows IS breakdown by order difficulty:

Exhibit 7

Comparison of aggregator (left) and modified aggregator (right) to control sample

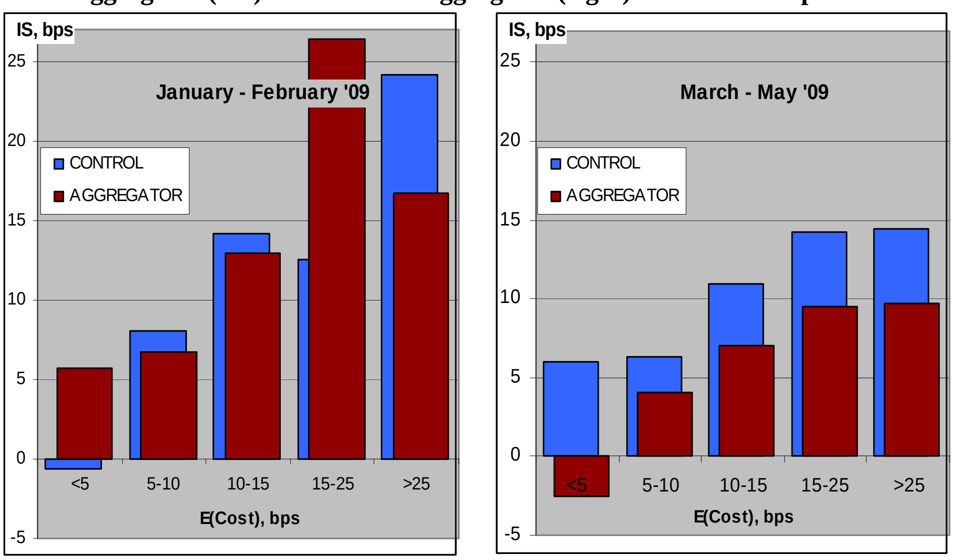

Implementation shortfall reduction is certainly encouraging. Still, we would like to better understand how the changes enhanced the aggregator's performance. Also there is room to improve robustness of our analysis: the IS measure is fairly noisy for larger trades and dollar weighing may skew the performance numbers. Also, the negative IS result for small trades in the modified aggregator may point to either limit price effects or selection bias – i.e. the aggregator being used more in relatively more favorable market conditions. The fill rates are comparable in all samples, allowing us to discard limit price effects. To address the possibility of selection bias we will need to break down costs below, using the methodology described in the first part of this paper to isolate adverse selection costs, opportunistic savings, alpha and market impact costs.

The section below compares our executions against the 20% PWP benchmark. We choose 20% to reflect our average execution rate (21-25% $-weighed, 17-18% simple weighted). We expect that market impact of "ideal" 20% strategy will on average be the same as ours. In this case out- or underperformance can be interpreted as an ability (or inability) of our strategies to execute at attractive price points and avoid falling prey to someone else's opportunistic trading.

There is a concern that for orders which were traded at a rate slower than 20%, the PWP window would only capture part of the order's impact. This may create a bias in favor of PWP. The bias should be minimal as most of the price impact is incurred in the early stage of the trade. Nonetheless, we analyzed our performance versus 10% PWP as well and found that all the results that we present below are still valid.

The standard deviation of performance versus PWP is fairly tight and the performance is not a function of the stock price (but it does depend on trade size). From this perspective it makes sense to replace the dollar-weighed performance metric with a more robust success rate, e.g. the probability of outperforming PWP by any magnitude, broken down by order difficulty:

Exhibit 8

Percentage of trades that led to opportunistic savings rather than adverse selection

		E(Cost) < 12 bps	E(Cost) > 12 bps
CONTROL	Jan-Feb	46%	53%
	Mar-May	45%	53%
AGGREGATOR	Jan-Feb	46%	57%
	Mar-May	**53%**	**60%**

We have noticed that small, "easy" orders fare a lot worse losing on average to PWP (and IVWAP) benchmarks. The majority of larger orders, however, outperform both these benchmarks. E(Cost) = 12 bps roughly draws the line between out- and underperformance in our samples. It's not that the same strategies execute difficult orders more intelligently. We still lose on average to high-frequency strategies, and it shows in small order performance (our success rate is 45-46%), but on the larger trades the adverse selection is masked by order's own impact, generally frontloaded execution, and trader controls.

The table above shows that the control strategies perform exactly the same across the two samples, and almost the same as the original aggregator: it is only marginally better for large orders. The modified aggregator, on the other hand, shows a very significant improvement in small orders, and further improves large order execution. The strategy ceased to be an easy target for adverse selection.

We have also found that the PWP performance depends on the price trend, or rather the average rate of price change in basis points per minute. A typical aggregator will get adversely selected in both favorable price trends (by filling an order too quickly), and in adverse trends (by executing too slowly). The table below confirms this behavior: in January-February the aggregator outperformed PWP only in non-trending market conditions. In the March – May sample, however, we have managed to maintain good performance in adverse trends as well. This can be attributed to the ability to go to the displayed markets and enforce a minimum participation rate while opportunistically searching for attractive liquidity. There is still room for improvement in favorably trending markets, though.

Exhibit 9

Occurrence of opportunistic savings versus prevailing market trend

		Highly adverse	Adverse	Neutral	Favorable	Highly Favorable
CONTROL	Jan-Feb	38%	39%	47%	51%	48%
	Mar-May	42%	46%	49%	48%	32%
AGGREGATOR	Jan-Feb	44%	39%	51%	48%	38%
	Mar-May	**57%**	**55%**	**60%**	46%	38%

Finally we calculate adverse selection (AS), opportunistic savings (OS), and the tracking error (TE) according to the formulas [2] – [4] in the earlier section. The aggregator shows an overall improvement of 3 bps versus a 0.4 bps improvement in the control set; this stems from a 37% reduction in adverse selection costs in the modified aggregator, while preserving essentially all the opportunistic savings. $-weighed PWP performance is also marginally better due to smaller improvements in large order performance.

Exhibit 10

Adverse selection and opportunistic savings

			#Orders	Traded, $ bn	$-weighted PWP, bps	AS | OS, bps	-TE = OS - AS, bps
CONTROL	Jan-Feb	AS	1335	0.93	(13.8)	(7.6)	
		OS	1101	0.95	18.9	8.5	**1.0**
	Mar-May	AS	1171	1.05	(13.7)	(7.4)	
		OS	997	1.11	19.2	8.8	**1.4**
AGGREGATOR	Jan-Feb	AS	640	0.49	(16.5)	(8.6)	
		OS	585	0.55	17.6	8.4	**(0.2)**
	Mar-May	AS	557	0.66	(12.3)	(5.4)	
		OS	712	0.71	14.7	8.2	**2.8**

While the reduction in adverse selection through modifications to the dark aggregator are remarkable, they do not completely explain the implementation shortfall average for the modified aggregator in Mar-May 2009, which was only 7.2 bps versus 16.0 bps in Jan-Feb 2009. To estimate selection bias, we use the decomposition of implementation shortfall into its components. Given a model to estimate market impact which we can take to be the expected impact of a 20% strategy with zero tracking error, $E(MI) = E(IS_{20\%})$, the implementation shortfall is

$$IS = E(IS_{20\%}) + \alpha + AS - OS$$

In the present case we have already seen that the expected cost is over-estimated probably because it uses lagging volatility values and the sample period followed a period of unusually high volatility. Therefore it is more accurate to derive an impact estimate from the observed shortfall from the control sample in the same Mar-May 09 period. The control sample shortfall was 11.5bps with a tracking error of -1.4bps (better than PWP), so the estimated market impact for that sample was 12.9bps. Using the ratio of expected costs (18/20) we can estimate the market impact for the trades executed in the modified aggregator to be 11.6bps. Given its tracking error of -2.8bps, the shortfall absent any alpha would be 8.8bps. We conclude that the modified aggregator sample was exposed to a favorable alpha bias of 1.6bps, leading to an observed average shortfall of 7.2bps. It should be noted that while the difference of 2.8 bps in IS is statistically significant at 95% confidence level, the 1.6 bps difference is not. The improvement attributable to the modifications in the aggregator therefore represent a reduction of 3/12.3 = 24% in the implementation shortfall.

This analysis illustrates the difficulties of dealing with real trade data, but also how the methodology described in this paper can help to unravel the mysteries in post-trade data by decomposing costs into its key components. The last illustration of this method tackles the puzzle of having negative shortfall observed for small orders (E(Cost) <5bps). For small orders in the aggregator the tracking error has changed from 4.8 bps (loss) in Jan-Feb to -3.4 (gain) in Mar-May. Implementation shortfall has improved accordingly. For the control set the tracking error has worsened from 0.9 to 7.2, and so is IS. (We use $-weighted TE to explain changes in $-weighed IS; success rate is obviously much less noisy metric as demonstrated by Exhibit 8.)

Implicit in the argument we used here for controlling for selection bias is the assumption that market impact (but not the PWP tracking error) is strategy-independent, for a fixed execution speed. We make this assumption primarily because it is very difficult to isolate what price effects are attributable to an algorithm's actions versus the actions of other market participants, given that one can never roll back the clock and replay the market without the presence of the algorithm. This topic is clearly deserving of further study.

Opportunistic algorithms have lower shortfalls than structured methods especially in volatile markets, yet they are more exposed to adverse selection. Is this a zero sum game?

The findings by Mittal and Wong indicate that algorithms with opportunistic strategies have lower implementation shortfalls than scheduled execution methods in volatile markets. This study considers only orders that were required to complete, but some of these orders inevitably get canceled as portfolio managers can choose to stop trading when the price moves more than they anticipated. Jeria and Sofianos show that the consideration of clean-up costs for canceled trades can significantly affect the conclusions as to whether an opportunistic trading method performs better or worse than a participation strategy. Recently, Altunata and Waelbroeck [Altunata and Waelbroeck (2008)] have shown that opportunist algorithms that lie on the Pareto-optimal frontier display both increased opportunistic savings and increased adverse selection when used in volatile markets. The relationship between adverse selection and volatility is reflected also here if we compare the results of the first section from the volatile markets in Nov-Dec 2008, where adverse selection was 11.2bps, to the results in the second section where we see adverse selection of the comparable aggregator at 8.6 in the January-February period, and for the control sample, 7.6 in Jan-Feb 09 to 7.4 in Mar-May 09. A similar relationship exists between opportunistic savings and volatility.

Although Pareto optimality implies that there is a tradeoff between impact and opportunity costs when running a single algorithm, the results in the previous section show that this tradeoff is not a zero-sum game; the 2/3 reduction in adverse selection in the engine more than compensates for a much smaller reduction in opportunistic savings. How is this possible? Both adverse selection and opportunistic savings result from deliberate actions: the former occurs when a passive order is executed (or not) depending on someone else’s short-term view; the latter is the result of an opportunistic algorithm exercising an option to trade aggressively at a good price point, or vice-versa, opting not to chase the stock on a short-term spike. In the context of dark aggregators, this suggests using dark aggregators tactically to access their liquidity at favorable price points but not reside there for an extended period of time.

As we stated in the introduction, adverse selection occurs when high frequency traders or arbitrageurs find that the stock is mispriced. An order becomes specifically vulnerable to adverse selection in the dark aggregator, because this is the first address an arbitrage trader will visit for price improvement and liquidity opportunities. After exhausting all the liquidity in the dark, the arbitrage traders will turn to the displayed market, causing the price to move towards fair value.

Much as arbitrage traders, institutional traders need to be aware of short-term alpha when deciding whether to use opportunistic or structured trading strategies. By using an opportunistic trading strategy at a time when the risk of adverse information is significant, the trader may be creating an opportunity for others through the implicit optionality in the order flow from any opportunistic algorithm.

If dark aggregators are subject to adverse selection, this also implies that the rate of executions in a dark aggregator reflects short-term alpha. The costs of adverse selection in the dark aggregator indicate that when execution rates rise in the dark pools, these are accompanied by strong trends in the opposite direction to trading. Conversely, the absence of liquidity in the dark aggregator can be a lead indicator of a market move in the same direction. This is a form of reverse information leakage: high frequency traders reveal the results of their alpha models through their activity in dark pools. The institutional trader can take advantage of this information in deciding the next step in a trading strategy.

Limit prices set by traders can also create adverse selection or opportunistic savings through their management of limit prices and aggression. If the price moves away and the trader is eventually forced to update the limit, the absence of fills in the early part of the trend can be very costly. Gomes and Waelbroeck [Gomes and Waelbroeck, 2009] have shown that on average the practice of manually managing tactical limit prices increases the risk exposure in a trade without a corresponding reduction in the average shortfall, when clean-up costs are counted (i.e., for non-discretionary orders that have a limited time horizon). On the other hand a policy of adopting more aggressive limits (higher limits for buys) at the beginning of a trade and controlling impact by using well-designed algorithms with low to moderate aggression levels leads to better execution results. Vice-versa, hands-on traders use occasional pull-backs during the second half of a trade (temporary market withdrawals or passive limits) and surgical strikes (aggressive short-term trades with tight limits) to capture attractive price opportunities after most of the position risk has been eliminated.

Based on the earlier observations [Altunata and Waelbroeck (2008)] and the results presented here, we recommend the following actions to avoid adverse selection:

- Avoid using passive strategies or manually managing tactical limits early in the trade, especially in high volatility markets when trading non-discretionary orders
- Use dark aggregators tactically with limited exposure times
- When using opportunistic algorithms or trading tools, it is important to control the participation rate. Ideally this can achieved pre-emptively through predictive analytics, but ex-post control measures like imposing minimum and maximum participation rates already lead to significant improvements.
- Exploit the alpha gleaned from the participation rate anomaly. On a rise in the dark fill rate, consider pulling back to passive strategies; when dark aggregators run dry consider engaging in the displayed markets to keep to a reasonable schedule.
- Use opportunistic algorithms or dark aggregators in situations with low adverse information risk

References

[Altunata and Waelbroeck, 2008] Serhan Altunata and Henri Waelbroeck, "Execution Styles and Trading Costs in the Pipeline Algorithm Switching Engine", Pipeline Financial research report, PIPE-2008-09-004, September 2008

[Farmer, Lillo, Waelbroeck 2007] Farmer, D., Lillo, F. and Waelbroeck, H. "Information Arbitrage and Market Impact", Pipeline Preprint PPLN-07-08-002 (2007)

[Farmer, Gerig, Lillo and Waelbroeck 2009] Farmer, D., Gerig, G., Lillo, F. and Waelbroeck, H. "Information Arbitrage and Market Impact", Pipeline Preprint PPLN-09-01-002 (2009)

[Glosten, Milgrom 1985] Glosten, L. and Milgrom, P., "Bid, ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders", J. Financial Economics 14, 71-100 (1985)

[Gomes and Waelbroeck, 2009] Carla Gomez and Henri Waelbroeck, "Effect of Trading Velocity and Limit Prices on Implementation Shortfall", Pipeline Financial research report, PIPE-2008-09-003, September 2008

[Jenkins 2001] Jenkins, Peter W., "Evaluating Cost-Measurement Tools", AIMR Conference Proceedings Vol. 2001, No. 2 : pp. 45-51, May 2001

[Jeria and Sofianos, 2009] David Jeria and George Sofianos, "Natural Adverse Selection: Theory and Evidence", Institutional Investor's "A Guide to Liquidity", January 2009

[Mittal and Wong, 2008] Hitesh Mittal, "Choosing the 'Right' Algorithm Can Minimize Trading Costs in Today's Volatile Markets", ITG Inc. research report, November 2008

[Stephens and Waelbroeck, 2009] Stephens, C. and Waelbroeck, H., "Algorithm switching: co-adaptation in the market ecology" to be published in Journal of Trading (2009)

Disclaimers: