

David Jeria
david.jeria@gs.com
NY: 917-343-6886

George Sofianos
george.sofianos@gs.com
NY: 212-902-9572

Related analysis:

Jeria & Sofianos, **Passive orders and natural adverse selection**, *Street Smart* Issue 33, September 4, 2008

Jeria & Sofianos, **Quantifying the SIGMA X crossing benefit: the sequel**, *Street Smart* Issue 32, July 8, 2008

Jeria & Sofianos, **Quantifying the SIGMA X crossing benefit**, *Street Smart* Issue 31, March 31, 2008

Goldman Sachs Equity Execution Strategies

This material has been prepared by Goldman Sachs Equity Execution Strategies and is not a product of Global Investment Research or Fixed Income Research. Important disclosures appear at the end of this material.

Passive in SIGMA X: quantifying the trade-offs

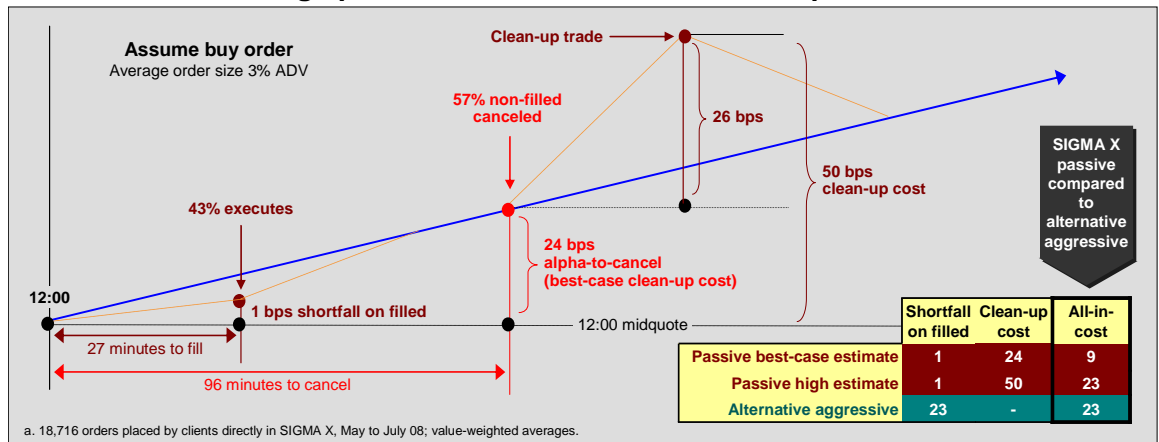
In *Street Smart* issue 33, we described how natural adverse selection creates a trade-off on passive orders between reduced market impact and increased short-term alpha loss.¹ The increased ST-alpha loss usually shows up in the high clean-up cost on the non-filled shares. To properly evaluate passive strategies, we proposed a framework that compares the all-in-cost (including the clean-up cost) of passive executions with the cost of an alternative aggressive strategy. In this issue of *Street Smart*, we use this framework to evaluate a sample of passive orders placed by clients directly in SIGMA X, the Goldman Sachs dark pool.

Exhibit 1 summarizes the behavior of the average passive SIGMA X order in our sample:

- The fill rate is 43 percent and the shortfall on filled shares is **only one basis point**.²
- The clean-up cost on non-filled shares, however, is high, ranging from 24 to 50 bps.³
- Adding the clean-up cost to the shortfall on filled shares, the all-in-cost ranges from 9 to 23 bps but still compares favorably with the cost of the alternative aggressive strategy.⁴

Our sample of passive SIGMA X orders has surprisingly high ST-alpha (41 bps) and high natural adverse selection spread between filled and non-filled shares (47 bps).⁵ **Despite high ST-alpha and natural adverse selection, the passive SIGMA X orders perform well relative to the alternative aggressive strategy.**

Exhibit 1. The average passive SIGMA X order in our sample^a



The averages in Exhibit 1 hide disparities across the three passive SIGMA X order types in our sample: midquote pegs (most popular), marketable limits and non-marketable limits. **Non-marketable limits perform best and substantially outperform the alternative strategy.** Nevertheless, the traders in our sample use non-marketable limits relatively infrequently. This suggests non-marketable limits are underutilized, most likely because of exaggerated concerns about natural adverse selection and gaming.

Midquote pegs do not perform as well as non-marketable limits. Our analysis suggests that **the average trader in our sample may benefit by canceling midquote pegs and switching to the alternative aggressive strategy more quickly.** Unlike non-marketable limits where the filled shares capture a generous spread reward that offsets natural adverse selection, midquote pegs typically execute at the midquote and offer little cushion against natural adverse selection.

We next discuss our findings in more detail.

Our sample of passive SIGMA X orders and natural adverse selection

Clients may place directly in SIGMA X marketable limits, non-marketable limits and pegged orders (pegged at midquote, bid or ask).⁶ Orders placed by clients directly in SIGMA X stay within SIGMA X until executed or canceled by the clients.⁷

Exhibit 2. Natural adverse selection: our sample of SIGMA X orders^a

	A RECEIVED				B FILLED		NON-FILLED		OVERALL	C Natural adverse selection spread ^c (bps)
	Orders (#)	Order size (shrs)	Order size (%ADV)	Arrival time	Fill rate (%)	Alpha-to-close ^b (bps)	Non-fill rate (%)	Alpha-to-close ^b (bps)	Alpha-to-close ^b (bps)	
All passive orders	18,716	13,431	3%	12:05	43%	14	57%	61	41	47
Midquote pegged orders	11,203	7,623	3%	11:59	34%	21	66%	52	41	31
Marketable limit orders	3,581	34,635	3%	12:09	50%	20	50%	75	48	55
Non-marketable limit orders	3,932	10,668	2%	12:04	39%	-19	61%	45	21	64

a. Passive orders placed by clients directly in SIGMA X, May to July 08; value-weighted averages.
 b. For buy orders, closing price minus midquote at arrival as % of midquote at arrival. Same calculation whether order filled or not.
 c. Alpha-to-close on non-filled orders minus alpha-to-close on filled orders.

Source: Goldman Sachs Equity Execution Strategies

“Toxic” clean-up trades
 Passive limits, high ST-alpha

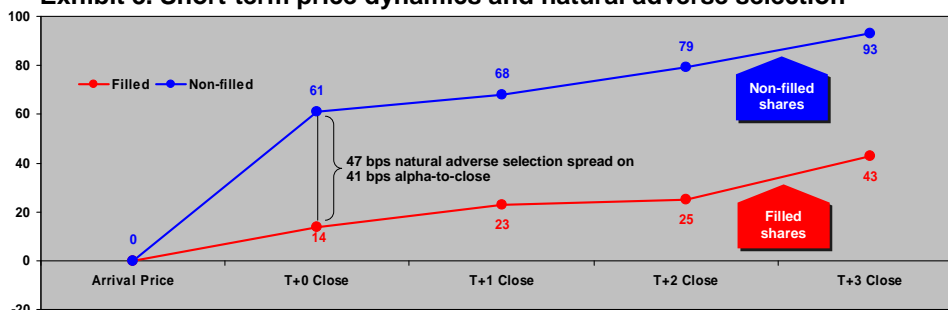
Our sample consists of orders placed directly in SIGMA X by hedge funds and traditional asset managers. Because we want to focus on the experience of asset managers when they directly place passive orders in dark pools, we exclude from our sample orders placed in SIGMA X by the Goldman Sachs algorithms, by other internal trading desks and by external broker-dealers.⁸ We also exclude all high-frequency flow irrespective of source (broker-dealer, electronic market maker or hedge fund).⁹

Exhibit 2 shows our final sample, 18,716 orders spread across all three SIGMA X order types:¹⁰ midquote pegs (11,203), marketable limits (3,581) and non-marketable limits (3,982).¹¹ Orders are evenly divided between buys and sells, and include large-cap, mid-cap and small-cap stocks.¹²

Our sample contains several surprises. We were surprised by the small number of orders hedge funds and traditional asset managers place directly in SIGMA X; over our three-month sample period they placed less than 20,000 orders. We were also surprised by the small size of these orders; on average, only three percent of ADV. Concerns about adverse selection and gaming are discouraging traders from placing large passive orders in SIGMA X. **Our results suggest that these concerns are not justified.**

Panel B in Exhibit 2 shows another surprise. In our sample, the overall ST-alpha measured by the alpha-to-close is high (41 bps).¹³ Even the patient non-marketable limits have high overall alpha-to-close (21 bps). The high ST-alpha is surprising because patient passive strategies are usually recommended for low ST-alpha flow. **Our results suggest that passive strategies may add value even when ST-alpha is high.**

Exhibit 3. Short-term price dynamics and natural adverse selection^a



a. 18,716 orders placed by clients directly in SIGMA X, May to July 08; value-weighted averages.

Source: Goldman Sachs Equity Execution Strategies

Our sample contains yet another surprise. Traditional asset managers are often reluctant to trade with hedge funds because of the presumption that hedge funds have more “toxic” (higher ST-alpha) flow. **The traditional asset managers in our sample, however, have similarly**

high ST-alpha as the hedge funds. This similarity highlights the difficulty of identifying and excluding from trading potential counterparties with “toxic” flow.

What is the natural adverse selection our sample orders are facing? In Exhibit 2, the alpha-to-close on filled orders is 14 bps while the alpha-to-close on non-filled orders is 61 bps. **The high 47 bps “spread” between the alpha-to-close on filled and non-filled orders is one possible measure of natural adverse selection.** If it was not for natural adverse selection, the average alpha-to-close on filled and non-filled shares would be the same and the natural adverse selection spread would be zero.

Exhibit 3 shows that the high alpha-to-close and natural adverse selection spread we observed in Exhibit 2 persists when we use T+1, T+2, or T+3 closing prices. Exhibit 3 also nicely illustrates natural adverse selection: the orders that should have filled quickly (alpha-to-close 61 bps) did not fill at all. **Given this high natural adverse selection, is a passive SIGMA X strategy justified?** To answer this question we next quantify the trade-off between natural adverse selection and reduced market impact.

Quantifying the natural adverse selection trade-off

Exhibit 4. Evaluation framework

Let:

- q = quantity received
- x = quantity filled
- z = quantity non-filled

$$q = x + z$$

- p = fill rate (x/q)
- $S_L(x)$ = shortfall on filled
- $S_C(z)$ = shortfall on non-filled (clean-up cost)
- $S_M(q)$ = shortfall on alternative aggressive strategy

All-in-cost of passive strategy:

$$p \cdot S_L(x) + (1-p) \cdot S_C(z)$$

Passive strategy is justified if its **all-in-cost** is less than cost of alternative aggressive strategy:

$$p \cdot S_L(x) + (1-p) \cdot S_C(z) < S_M(q)$$

Post-trade, the only numbers we know with certainty are the fill rates (p) and the actual shortfall on the filled quantity (S_L); we must estimate the cost of the alternative aggressive strategy (S_M) and the clean-up cost (S_C).

Exhibit 4 summarizes our framework for quantifying the natural adverse selection trade-off. We first estimate the all-in-cost of the passive strategy including the clean-up cost and then compare it with the estimated cost of a benchmark alternative aggressive strategy. If the all-in-cost of the passive strategy is less than the cost of the alternative strategy, the passive strategy is working.

Exhibit 5 shows our estimation of the cost of the alternative strategy. The shortfall on the alternative strategy has two components:

- **Market impact.** We estimate market impact using the Goldman Sachs Shortfall Model (GSSM). We assume the alternative strategy begins executing immediately on order arrival, at 20 percent participation, and the order must fill completely. The estimated average impact of the alternative strategy is 19 bps.
- **Execution horizon (EH) alpha loss.** As prices change over the execution horizon, the alternative strategy may also incur alpha loss. We estimate the EH-alpha loss by pro-rata allocating the actual alpha-to-close on each order in our sample in proportion to the order’s expected half life (34 minutes).¹⁴ The EH-alpha loss on the alternative strategy is 3 bps.

Adding the impact and alpha-loss, the shortfall on the alternative strategy is 23 bps.

We next estimate the all-in-cost of the passive executions in our sample. Exhibit 6 shows the actual shortfall on the filled shares. The average fill rate is 43 percent and the actual shortfall on the fills is one basis point. With 43 percent fill rate, just looking at the shortfall on filled shares is misleading. **To properly evaluate executions with low fill rates we must also estimate the clean-up cost.**

Exhibit 5. The alternative aggressive strategy^a

	ALTERNATIVE STRATEGY			
	Expected impact ^b (bps)	Expected half-life (min)	Expected EH-alpha loss ^c (bps)	Expected all-in-cost ^d (bps)
All passive orders	20	34	3	23
Midquote pegged orders	21	37	4	25
Marketable limit orders	19	35	4	23
Non-marketable limit orders	17	26	1	18

a. Passive orders clients place directly into SIGMA X, May to July 08; value-weighted averages.
b. From Goldman Sachs Shortfall Model (GSSM), execution start at arrival time, 20% participation.
c. Alpha-to-close allocated in proportion to expected half life (strike time to completion divided by two).
d. Expected impact plus expected EH-alpha loss

Source: Goldman Sachs Equity Execution Strategies

Exhibit 6. Shortfall on filled component^a

	PASSIVE STRATEGY		
	Filled quantity		
	Fill rate (%)	Actual shortfall ^b (bps)	Execution half-life ^c (minutes)
All passive orders	43%	1	27
Midquote pegged orders	34%	13	33
Marketable limit orders	50%	3	23
Non-marketable limit orders	39%	-30	36

a. Passive orders clients place directly into SIGMA X, May to July 08; value-weighted averages.
b. For buys, execution price minus midquote at order arrival as % of midquote.
c. For each order with multiple child orders we first calculate exec time for each child (arrival to exec) and then calculate the value-weighted mean across all child execution times.

Source: Goldman Sachs Equity Execution Strategies

Estimating the cost of the clean-up trade

Estimating the clean-up cost is difficult because a clean-up trade may never occur and even when a clean-up trade does occur it may be hard to identify. Suppose XYZ trades at \$20.10

Exhibit 7. The cost of the clean-up trade^a

	PASSIVE STRATEGY					
	The clean-up trade					
	Non-fill rate (%)	Actual time to cancel ^b (min)	Actual alpha-to-cancel ^c (bps)	Expected impact ^d (bps)	Expected EH-alpha loss ^e (bps)	Expected clean-up cost ^f (bps)
A Clean-up trades execute at midquote on cancel time (best case)						
All passive orders	57%	96	24	0	0	24
Midquote pegged orders	66%	92	18	0	0	17
Marketable limit orders	50%	86	33	0	0	33
Non-marketable limit orders	61%	128	13	0	0	13
B Clean-up trades begin executing at cancel time at 20% participation						
All passive orders	57%	96	24	19	7	50
Midquote pegged orders	66%	92	18	22	8	48
Marketable limit orders	50%	86	33	19	7	59
Non-marketable limit orders	61%	128	13	17	5	35

a. Passive orders clients place directly into SIGMA X, May to July 08; value-weighted averages.
b. Cancel time minus order arrival time.
c. For buys, midquote at cancel time minus midquote at arrival as % of midquote at arrival.
d. Bottom panel: from Goldman Sachs Shortfall Model (GSSM), execution start at order arrival, 20% participation.
e. Bottom panel: half of alpha from order arrival to end of 20% participation or to close whichever comes first.
f. Alpha-to-cancel plus expected impact plus expected EH-alpha loss.

Source: Goldman Sachs Equity Execution Strategies

and AlphaMax places a limit order in SIGMA X to buy XYZ at \$20.00 or better. The price of XYZ rises 90¢ to \$21.00 and the limit is not triggered. AlphaMax cancels the SIGMA X order and decides not to buy XYZ. Even though no clean-up trade occurs, AlphaMax's non-filled order still has a cost: the missed opportunity to capture the 90¢ price move net of the cost of the clean-up trade had it actually occurred. **To properly evaluate passive executions, therefore, we must assign a hypothetical clean-up cost even if a clean up trade did not occur and compare with the missed alpha opportunity.**

Now suppose when XYZ hits \$21, AlphaMax again cancels the order in SIGMA X but now buys XYZ at the market through another venue. In this case a clean-up trade did occur but is not in our data. AlphaMax may also have a

hard time tracking clean-up trades and tying them to the original execution attempts.

For all these reasons, instead of trying to identify the actual clean-up trades, we estimate the cost of hypothetical clean-up trades. **We anchor our estimates to reality by using the actual alpha-to-cancel on non-filled orders.** In our data we capture the cancel time a client cancels a non-filled order. We use this cancel time to calculate the alpha-to-cancel, the price move from arrival to cancel.¹⁵ In Exhibit 7, the average time-to-cancel on the non-filled shares is 96 minutes and the average alpha-to-cancel is 24 bps. With natural adverse selection working against the order, by the time the average client in our sample canceled her non-filled orders, the price moved 24 bps away from her. **The alpha-to-cancel is a real, sunk cost of the clean-up trade and we can accurately measure it.**

We calculate two sets of clean-up-cost estimates. In Exhibit 7 Panel A, we assume the client executes the non-filled quantity at the midquote immediately on canceling the order in SIGMA X. In this best-case scenario the clean-up-cost is just the alpha-to-cancel (24 bps).¹⁶ In Panel B, we estimate the clean-up cost assuming 20 percent participation. The 20 percent clean-up trade mirrors the alternative aggressive strategy (Exhibit 5) but now calculated on the residual quantity and beginning at cancel time. The 20 percent clean-up trade has a market impact component which we estimate from GSSM (19 bps) and an EH-alpha loss component which we estimate by pro-rata allocating the actual alpha-to-close (7 bps).¹⁷ Summing up the alpha-to-cancel, market impact and EH-alpha loss, the total cost of the clean-up trade is 50 bps.

The wide range between the two clean-up estimates (24 to 50 bps), underscores the difficulty of estimating the clean-up cost. Most probably, the "truth" lies somewhere in between these two estimates.

Putting it all together: the verdict on SIGMA X passive executions

Exhibit 8 pulls together the statistics we use to evaluate passive orders.¹⁸ **The all-in-cost of the passive strategy is the weighted average of actual shortfall on filled shares and clean-up cost on non-filled shares, using fill rates as weights.** We generate two all-in-cost estimates: a "low" estimate using the best-case clean-up cost and a "high" estimate using the 20 percent participation clean-up cost.

Consider the low estimate first. Shortfall on the 43 percent filled is one basis point, the clean-up cost on the 57 percent non-filled is 24 bps, and the all-in-cost is 9 bps.¹⁹ The estimated

cost of the alternative aggressive strategy is 23 bps, so the low estimate for the all-in-cost of passive SIGMA X executions is substantially better than the alternative strategy. For the high estimate, the clean-up cost is 50 bps, so the all-in-cost of the passive strategy is 23 bps, the same as the cost of the alternative strategy. **Even our high estimate of the all-in-cost of passive SIGMA X executions compares favorably with the cost of the alternative strategy.**

Of the three SIGMA X order types, non-marketable limits perform best even though they face the highest natural adverse selection (64 bps, Exhibit 2). They perform best because the filled shares (39 percent) execute 30 bps better than the arrival midquote. This 30 bps spread capture creates a comfortable cushion to offset high natural adverse selection and the clean-up cost on the non-filled shares (13 to 35 bps).²⁰ **The all-in-cost on non-marketable limits, therefore, is only 6 bps (high estimate) compared to 18 bps for the alternative strategy.**

Midquote pegs and marketable limits perform less well but still compare favorably with the alternative strategy. On midquote pegs, the fill rate is 34 percent, the shortfall on fills is 13 bps and the clean-up cost is 47 bps (high estimate). The all-in-cost of midquote pegs is 27 bps (high estimate) compared to 25 bps for the alternative strategy. Midquote pegs are expensive because the 13 bps shortfall on the filled shares, unlike the 30 bps spread capture on non-marketable limits, does not provide a cushion to offset the high clean-up cost.

Exhibit 8. Putting it all together^a

	RECEIVED			PASSIVE STRATEGY					ALTERNATIVE STRATEGY		
	Orders (#)	Order size (shrs)	Arrival time	Filled		Non-fill clean-up cost		All-in-cost		Expected all-in-cost ^e (bps)	
				Fill rate (%)	Actual shortfall ^b (bps)	Non-fill rate (%)	Expected clean-up cost ^c (bps)	Expected all-in-cost ^d (bps)			
All passive orders	18,716	13,431	12:05	43%	1	57%	24	50	9	23	23
1 Midquote pegged orders	11,203	7,623	11:59	34%	13	66%	17	48	12	27	25
2 Marketable limit orders	3,581	34,635	12:09	50%	3	50%	33	59	15	26	23
3 Non-marketable limit orders	3,932	10,668	12:04	39%	-30	61%	13	35	-12	6	18

a. Passive orders clients place directly into SIGMA X, May to July 08; value-weighted averages.

b. For buys, execution price minus midquote at order arrival as percent of midquote.

c. Actual alpha-to-cancel plus 50% of expected impact from Goldman Sachs Shortfall Model (GSSM) at 20% participation plus expected EH alpha.

d. Weighted average of actual shortfall on filled and expected clean up cost (the mid estimate) using the fill rates as weights.

e. Expected impact from GSSM at 20% participation rate plus expected EH-alpha loss.

Source: Goldman Sachs Equity Execution Strategies

The filled non-marketable limits execute 30 bps better and the midquote pegs 13 bps worse than the arrival midquote. This large discrepancy highlights the difference between these two passive strategies. For non-marketable limits, natural adverse selection only affects the fill rate. When the order fills, the trader controls the execution price: at or better than the specified limit price. For midquote pegs, however, the trader controls neither the fill rate nor the execution price, so natural adverse selection affects both. Because of natural adverse selection, midquote executions will on average occur at worse midquotes than the arrival midquote. In our sample, the 13 bps shortfall relative to the arrival midquote is caused by an adverse drift in the midquote from arrival to execution. **Our results suggest that the average trader in our sample may benefit by canceling midquote pegs and switching to the alternative aggressive strategy faster.**²¹ We discuss midquote pegs further in the Appendix.

On marketable limits, the fill rate is 50 percent, the shortfall on fills is 3 bps, and the clean-up cost is 50 bps (high estimate). The all-in-cost of midquote pegs is 26 bps (high estimate) compared to 23 bps for the alternative strategy. Marketable limits are expensive because of the high clean-up cost. Even in the best case scenario the clean-up cost (alpha-to-cancel) is 33 bps. **This again suggests that the average trader in our sample may be waiting too long before canceling her marketable limits.**²²

The averages in Exhibit 8 hide a range of experiences across the different clients in our sample. Some clients' passive orders perform better and some perform worse than the average. **We provide the analytical tools to help clients fine-tune and improve their passive executions.**

Concluding comments

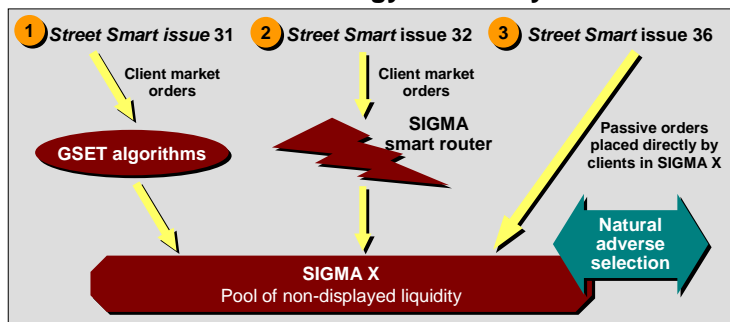
This report concludes a *Street Smart* trilogy in which we examine the three ways clients may benefit from SIGMA X, the Goldman Sachs pool of non-displayed liquidity (see Exhibit 9):

- 1. Using Goldman Sachs algorithms.** The child orders generated by Goldman Sachs algorithms interact with SIGMA X liquidity. In *Street Smart* Issue 31, we showed that for the VWAP algorithm, SIGMA X crossing reduces shortfall by 12 percent.
- 2. Using the Goldman Sachs SIGMA smart router.** When a client routes a marketable order through SIGMA, the smart router first exposes the order to SIGMA X liquidity. In *Street Smart* Issue 32, we showed that SIGMA X crossing reduces execution shortfall on orders that outsize quoted depth. On orders three times quoted depth, for example, SIGMA X crossing reduces shortfall by 25 percent.
- 3. Placing passive orders directly in SIGMA X.** In this issue of *Street Smart*, we show that passive SIGMA X executions compare favorably with the alternative aggressive strategy, despite the inevitable natural adverse selection that passive orders face.

The popularity and proliferation of dark pools, makes it increasingly important for traders to

properly evaluate the execution quality they receive on passive orders placed in these venues. Unfortunately, dark pools typically only report the execution quality on filled shares. Fill rates are rarely reported and no attempt is made to evaluate the clean-up cost on non-filled shares. Our analysis shows that this is misleading. In our sample, the shortfall on filled shares is only one basis point, but including the clean-up cost the all-in-cost is 23 bps. Given the low fill rates on passive orders in dark pools, ignoring the clean-up cost is equivalent to staging Hamlet without the prince! **We hope our all-in-cost**

Exhibit 9. Street Smart trilogy: three ways to SIGMA X



approach will become the standard on how to properly evaluate and optimize passive executions.

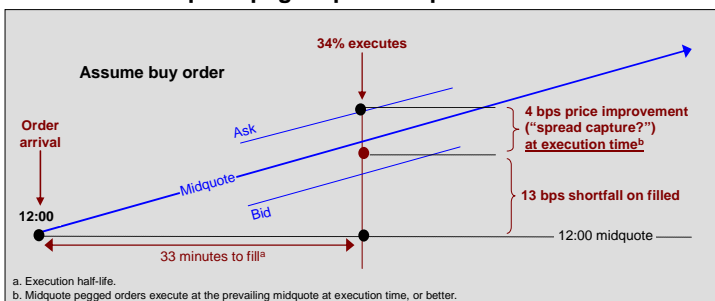
The proliferation of dark pools also makes it increasingly important for traders to take full advantage of this alternative source of liquidity. Our empirical analysis of passive SIGMA X executions suggests that asset managers benefit by placing non-marketable limit orders directly in dark pools. Traders, however, are not taking full advantage of this opportunity to add value, most likely because of concerns about adverse selection and gaming. **We hope that the empirical evidence we present in this report will re-assure traders.**

APPENDIX: Midquote pegs and the “spread capture” confusion

QUIZ: Why is the average all-in-cost of midquote pegs inevitably positive?

Exhibit A1 provides more execution detail on midquote pegs. The filled midquote pegs in our sample

Exhibit A1. Midquote pegs: spread capture but at what cost?



show 4 bps price improvement relative to execution-time best ask (buys) or bid (sells). This price improvement is confusingly also referred to as “spread capture.” This is confusing because the spread is captured by at-the-quote limits and not by the midquote pegs. At-the-quote limits buy at the bid and sell at the ask and so capture the bid-ask spread (see sidebar). The bid-ask spread is the roundtrip reward to liquidity providers that place at-the-quote limits. Part of this reward compensates liquidity providers for the natural adverse selection they face.²³ Midquote pegs, however, buy at midquote and sell at midquote, so the roundtrip spread capture is really zero. Because midquote pegs inevitably face natural adverse selection and because spread capture is zero, their average all-in-cost will inevitably be positive. Midquote pegs do not provide a spread capture “cushion” to offset the ST-alpha loss that will arise

from natural adverse selection.

This observation does not necessarily mean that midquote pegs are a bad execution strategy. The all-in-cost of midquote pegs is inevitably positive but it may still be lower than the cost of the alternative strategy. The lack of spread capture, however, makes it even more important for traders to evaluate midquote pegs by looking at the all-in-cost and not just price improvement at execution time.

Back in Exhibit A1, focusing on the 34 percent filled shares, the average time to execution is 33 minutes and execution shortfall is 13 bps. As usual, we measure execution shortfall for buy orders as execution price minus midquote at order arrival. Since midquote pegs by design execute at the midquote and have no market impact,²⁴ the 13 bps shortfall measures the pure alpha-loss from arrival to execution. **The average midquote peg in our sample, while waiting 33 minutes for a midquote match, misses the market by 13 bps.** Yes, the order eventually executes at the midquote and receives 4 bps price improvement but the midquote at execution time is 13 bps worse than at arrival.

The same argument applies to all passive midquote executions and not just midquote pegs. On crossing networks that match counterparties at the midquote, for example, these midquote crosses do not compensate passive orders for natural adverse selection.

Who is capturing the spread?

- Market orders that buy at the ask and sell at the bid “pay the spread”
- Limit orders that buy at the bid and sell at the ask “capture the spread”
- Midquote pegs that buy at the midquote and sell at the midquote neither pay nor capture the spread

¹ Natural adverse selection is the natural tendency for passive orders to fill quickly when they should fill slowly and fill slowly (or not at all) when they should fill quickly. We define passive orders broadly to include any order where the trader does not fully control the timing of execution. For a full discussion of natural adverse selection see Jeria & Sofianos, “Passive orders and natural adverse selection,” *Street Smart* Issue 33, September 4, 2008.

² For buy orders, we define execution shortfall as the execution price minus the midquote at order arrival as percent of the midquote at order arrival, in basis points. For sell orders execution shortfall is the midquote at order arrival minus the execution price as percent of midquote at order arrival.

³ We discuss the clean-up cost (the cost of the clean-up trade) in detail later in the report, see Exhibit 7.

⁴ The all-in-cost is a weighted average of the shortfall on the filled shares and the clean-up cost on the non-filled shares with the fill rates as the weight; we discuss the alternative strategy later in the report, see Exhibit 5.

⁵ Natural adverse selection spread is the difference between ST-alpha on filled and non-filled shares, see Exhibit 3.

⁶ On SIGMA X pegged orders, clients may also specify an absolute limit price (e.g. buy at midquote at \$20 or lower).

⁷ Day orders if not filled will cancel automatically at the end of the day. Unlike the SIGMA smart router, SIGMA X does not pro-actively route orders to external venues. For a discussion of the SIGMA smart router see Jeria & Sofianos “Quantifying the SIGMA X crossing benefit: the sequel,” *Street Smart* Issue 32, July 8, 2008.

⁸ For an analysis of how SIGMA X improves the execution quality of Goldman Sachs algorithms, see Jeria & Sofianos “Quantifying the SIGMA X crossing benefit,” *Street Smart* Issue 31, March 31, 2008.

⁹ To make sure we exclude all high frequency flow, we also drop all orders with completion or cancel time less than five minutes. Our qualitative results are not sensitive to the five minute cut-off.

¹⁰ As usual, we filter our sample for extreme orders. For example, we drop orders in Pink Sheet stocks, orders in low-priced stocks (less than \$1) and orders in high-priced stocks (more than \$150). In our final sample 2,136 orders did not fill at all, 9,793 orders completely filled and 6,787 orders partially filled.

¹¹ On pegged orders, the data had few bid and ask pegs, so we focus on midquote pegs, the most popular peg type.

¹² Buy orders are 48 percent of dollar value received. Large-cap stocks (> \$7.5bn) are 55 percent of value received, mid-cap 35 percent and small cap (< \$1 bn) 10 percent.

¹³ “Overall” alpha-to-close is the alpha-to-close irrespective of whether the order executed or not. On each buy order, irrespective of whether it filled or not, overall alpha to close is the closing price minus the midquote at order arrival as percent of the arrival midquote.

¹⁴ For details on the pro-rata allocation method of estimating EH-alpha see Abrokwah & Sofianos, “Shortfall Surprises,” *Journal of Trading*, Summer 2007 (in particular Exhibit 4). We get estimates of the execution half life from GSSM.

Average execution half-life in our sample is 34 minutes and average order arrival time is 12:05 so the allocation factor is about 1/8 and the overall alpha-to-close being allocated is 41 bps (these are value-weighted averages so doing the calculation using these averages does not give the 3 bps estimate in Exhibit 5).

¹⁵ For buy orders we define alpha-to-cancel as the midquote at cancel time minus the midquote at order arrival as percent of the arrival midquote. For sell orders we define alpha-to-cancel as the midquote at arrival minus the midquote at cancel time as percent of the arrival midquote. If the cancel is at the close we use the closing price.

¹⁶ Although this best-case scenario does not frequently take place, occasionally clients do cancel orders because they found midquote matches elsewhere.

¹⁷ If an actual clean-up trade did occur and the impact of that trade affected the stock price, it is possible that our EH-alpha loss estimate double counts market impact and overstates the cost. We are careful to use the closing price to estimate the EH-alpha loss because the closing price is less likely to be affected by the impact of the clean-up trade especially since the orders in our sample are relatively small. Some residual double-counting may remain so the high estimate is best viewed as an *upper bound* to the true clean-up cost.

¹⁸ The first row in Exhibit 8 shows the averages behind Exhibit 1.

¹⁹ Because we report *value-weighted averages* the all-in-cost is not equal to 43 percent times one basis point plus 57 percent times 24 bps. We calculate the all-in-cost as the aggregate dollar shortfall on filled shares plus the aggregate dollar clean-up cost on non-filled shares divided by the dollar value of the aggregate shares submitted for execution.

²⁰ We use the term “spread capture” to indicate negative execution shortfall. So our definition of spread capture is relative to the midquote at order arrival (for buy orders, midquote at arrival minus execution price).

²¹ But canceling more quickly will reduce the fill rate.

²² Canceling more quickly will reduce the alpha-to-cancel but also reduce the fill rate, so the trader should experiment and fine tune the strategy.

²³ For the role of natural adverse selection (or “information asymmetries”) in the theory of the bid-ask spread see Walter Bagehot (pseudo.) “The Only Game in Town,” *Financial Analysts Journal*, March 1971, Thomas Ho & Hans Stoll “Optimal Dealer Pricing and Return Uncertainty,” *Journal of Financial Economics*, March 1981, Thomas Copeland and Dan Galai “Information Effects on the Bid-Ask Spread,” *Journal of Finance*, December 1983.

²⁴ In practice, midquote pegs may receive price improvement, for example buy below the midquote.

This material has been prepared by personnel in Securities Division Sales & Trading area of one or more affiliates of The Goldman Sachs Group, Inc. (“Goldman Sachs”) and is not the product of Global Investment Research or Fixed Income Research. It is not a research report and is not intended as such.

Non-Reliance and Risk Disclosure: This material should not be construed as an offer to sell or the solicitation of an offer to buy any security in any jurisdiction where such an offer or solicitation would be illegal. We are not soliciting any action based on this material. It is for the general information of our clients. It does not constitute a recommendation or take into account the particular investment objectives, financial conditions, or needs of individual clients. Before acting on any advice or recommendation in this material, you should consider whether it is suitable for your particular circumstances and, if necessary, seek professional advice. Certain transactions - including those involving futures, options, equity swaps, and other derivatives as well as non-investment-grade securities, foreign-denominated securities and securities, such as ADRs, whose value is influenced by foreign currencies - give rise to substantial risk and may not be available to or suitable for all investors. This material is not for distribution to private customers, as that term is defined under the rules of the Financial Services Authority in the United Kingdom; and any investments, including derivatives, mentioned in this material will not be made available by us to any such private customer. The material is based on information that we consider reliable, but we do not represent that it is accurate, complete or up to date, and it should not be relied on as such. Opinions expressed are our current opinions as of the date appearing on this material and only represent the views of the author and not those of Goldman Sachs, unless otherwise expressly noted.

Legal Entities Disseminating this Material: This material is disseminated in Australia by Goldman Sachs JBWere Pty Ltd (ABN 21 006 797 897) on behalf of Goldman Sachs; in Canada by Goldman Sachs Canada Inc. regarding Canadian equities and jointly by Goldman, Sachs & Co. and Goldman Sachs Execution & Clearing, L.P. regarding all other materials; in Hong Kong by Goldman Sachs (Asia) L.L.C.; in Japan by Goldman Sachs Japan Co., Ltd.; in the Republic of Korea by Goldman Sachs (Asia) L.L.C., Seoul Branch; in New Zealand by Goldman Sachs JBWere (NZ) Limited on behalf of Goldman Sachs; in Singapore by Goldman Sachs (Singapore) Pte. (Company Number: 198602165W); in Europe by Goldman Sachs International (unless stated otherwise); in France by Goldman Sachs Paris Inc. et Cie and/or Goldman Sachs International; in Germany by Goldman Sachs International and/or Goldman, Sachs & Co. oHG; in Brazil by Goldman Sachs do Brasil Banco Múltiplo S.A.; and in the United States of America, unless otherwise noted, by each of Goldman, Sachs & Co. and Goldman Sachs Execution & Clearing, L.P. (both of which are members NASD, NYSE and SIPC). Goldman Sachs International, which is authorized and regulated by the Financial Services Authority, has approved this material in connection with its distribution in the United Kingdom and European Union. Unless governing law permits otherwise, you must contact a Goldman Sachs entity in your home jurisdiction if you want to use our services in effecting a transaction in the securities mentioned in this material.

Reproduction and Re-Distribution: No part of this material may be (i) copied, photocopied or duplicated in any form by any means or (ii) redistributed without our prior written consent.