ACE Discretionary Versus ACE Non-Discretionary: A Case Study Using ITG Algorithms Execution Data

MILAN BORKOVEC, DAVID HAN, AND HANS G. HEIDLE

July 18, 2007

Abstract

Starting with version 2.3, the ITG's Agency Cost Estimator (ACE®) provides two different cost estimates: ACE Discretionary and ACE Non-Discretionary. The difference between the two cost estimates is that the calibration is based on different samples of realized costs. ACE Discretionary is calibrated using the entire sample in the ITG Peer Group Database. For ACE Non-Discretionary, we use a subsample of the ITG Peer Group Database. In particular, opportunistic trading is identified with sophisticated heuristics and then excluded from the sample. This paper uses ITG Algorithms℠ execution data to validate the identification and filtration of opportunistic trading by comparing various subsamples of ITG Algorithms execution data with their respective ACE Discretionary and ACE Non-Discretionary cost estimates. Overall, the analysis shows that the filtration algorithm works well and ACE Discretionary and Non-Discretionary provide realistic and accurate cost estimates.

MILAN BORKOVEC is Head of the Financial Engineering Group at Investment Technology Group, Inc., 44 Farnsworth Street, Boston MA 02210, Tel: (617) 692-6733; e-mail: milan.borkovec@itg.com

DAVID HAN is a Research Analyst in Financial Engineering's Pre-Trade Analytics Group at Investment Technology Group, Inc., 44 Farnsworth Street, Boston MA 02210, Tel: (617) 692-6860; e-mail: david.han@itg.com

HANS G. HEIDLE is Manager of Financial Engineering's Pre-Trade Analytics Group at Investment Technology Group, Inc., 44 Farnsworth Street, Boston MA 02210, Tel: (617) 692-6828; e-mail: hans.heidle@itg.com

The authors would like to thank Yossi Brandes for his support and comments and suggestions.

1. OVERVIEW

Fundamentally, an order represents a demand for liquidity. This demand might be either absolute/deterministic or exhibit a discretionary component that can result in no or partial order execution if market conditions are unfavorable. The more discretion one has, the more likely it is to execute the order (or parts of the order) at a favorable price. This is typically done using fixed limit orders, delayed trading commencement, security substitution, and opportunistic crossing/block trading. The net effect of this "opportunistic" trading style is that trade performance for executed shares is very often above what is achievable by a market participant obligated to demand a fixed amount of liquidity.

This disparity in performance is reflected in the ITG Peer Group Database and, therefore, has the potential to affect price impact functions. For those seeking a cost estimate reflecting a determined, non-opportunistic trading style, price impact estimates should be based on a dataset that is free from executions of orders with discretion. For those seeking cost estimates that reflect what market participants in aggregate pay, a dataset including all orders is appropriate. The suitability of each of these estimates is guided by the nature of the orders to be benchmarked, and will vary by institution and within an institution, by manager or investment style.

To accommodate the need for two benchmarks for identical orders (besides the amount of discretion), starting with ACE/2.3, ACE/2 has the ability to provide two different cost estimates – one based on orders that have been fully executed no matter how the market conditions were and another based on all executed orders.[1] From a pre-trade perspective, the *ACE Non-Discretionary* estimate is highly suitable for vetting trading strategies and determining the feasibility of executing an order in its entirety. The more general *ACE Discretionary* cost estimate provides a number that is suitable for comparing incurred transaction costs with what other market participants experience. Systematically under- (or over-) performing compared to this number might suggest a trend in an institution's competitiveness. From a post-trade perspective, the choice of price impact models generally should be guided by the prevailing nature of the order. For example, orders that require

[1] For a detailed description of ACE, please see the white paper "ITG ACE – AGENCY COST ESTIMATOR" (2007), ITG INC .

immediate and continuous trading until completion should be compared against a cost estimate derived from a price impact model that reflects determined, non-opportunistic trading (ACE Non-Discretionary). However, an exception to this might be if the impetus for order creation frequently results from an observation of favorable market conditions or if orders are often not fully executed.

The associated price impact coefficients for ACE Non-Discretionary and ACE Discretionary are derived from different subsets of the same ITG Peer Group Database. For the ACE Discretionary model, the entire database minus orders eliminated by outlier filters is included in the calibration process. For the ACE Non-Discretionary model, a sophisticated set of heuristics is used to eliminate database participants that exhibit opportunistic trading. The heuristics focus on identifying participants whose orders do not meet minimum transaction cost requirements with respect to order size. More precisely, all orders of clients who have unusually low average transaction costs for a given exchange, liquidity group, and order size segment are filtered out. Liquidity groups are defined based on the deciles of the average daily dollar volume distribution from all stocks. Order size segments are defined as 0-1%, 1-5%, 5-10%, 10-25%, 25-50%, and >50% of the 21-day median daily share volume (MDV). The grouping is justified by the fact that different accounts or portfolio managers can trade very differently within the same firm. Actual average costs of a client are considered to be abnormally low (signaling opportunistic trading) if they are lower than a cutoff for the specific segment. The cutoff is determined by two thresholds:

a) based on a certain cutoff that equals the average half-spread for all orders multiplied by a certain factor for the given segment (e.g., the factor is one for order sizes around 15% of MDV),
b) based on the average realized cost of all market participants.

If the average trading costs of a client are less than both of the two thresholds described by a) and b) above, the client's trading style for this segment is classified as opportunistic and is filtered out for ACE Non-Discretionary.

This heuristic approach is motivated by the fact that some market participants have extremely low average transaction costs (even for (very) large order sizes), forcing the calibrated ACE cost estimates to be very low. On average, half of the ITG Peer Group Database clients, which corresponds to approximately 40 to 50 clients, are filtered out. Note that cutoff a) is typically binding for small and medium size orders while cutoff b) is binding mostly for larger order sizes.

The goal of this paper is to validate the above described heuristics for identifying opportunistic trading by using ITG Algorithms execution data. To accomplish that, we compare actual trading costs, ACE cost estimates, and the cutoff a) for ITG Algorithms data.

2. DATA

The ITG Algorithms data sample used for our analysis contains 1,119,460 orders over the period from October 2006 to March 2007 and consists of mostly small to medium size orders which motivates why we neglect cutoff b) in this study. Five different ITG Algorithms are used in this study: Dark, Active, Volume Participation (VP), Volume-Weighted Average Price (VWAP), and Implementation Shortfall (IS). The number of observations for each ITG Algorithm is reported in Table 1. The various algorithms have different degrees of discretion built in as described in the following.

We aggregate trades by trade date, moniker id, stock id, and side (*Cluster Set 1*). The entire sequence (*cluster*) of trades is treated as the basic unit of analysis in order to determine price impact and execution cost of institutional trading. In particular, we define a "buy cluster" to include the company's successive purchases of the stock. "Sell clusters" are defined similarly. The construction of clusters as proxy for ex-ante orders may be problematic in some situations, for example,

(i) favorable price movements subsequent to a trade may generate further orders and, consequently, trades of the same stock, creating a proximate sequence of trades; or

(ii) costs of execution risk might be underestimated since all clusters contain realized executions only, therefore, ex-ante orders are orders that have been fully completed. However, opportunity costs of unexecuted orders may be large creating a problem in comparing pre-trade cost estimates and post-trade costs.

We hope that (i) is not a problem since each trade was executed through an algorithm (although it is possible that clients can feed additional orders to the algorithm during the day). Item (ii) is definitely a problem. First, not all of the above-mentioned ITG Algorithms assume 100% order completion as a necessity. In fact, VWAP and IS are the only ones that do most of the times and executed volume in Dark can be even zero. Second, like other algorithms in the industry, ITG

Algorithms allow the external clients to specify price thresholds above or below which the orders are not being executed (so called *Algorithmic Limit Orders*). Algorithmic Limit Orders should not be mistaken with “regular” limit orders. Algorithmic Limit Orders refers to the fact that a client can specify price limits for the algorithm that restricts trading outside of these limits. The algorithm can still use any combination of market and limit orders for individual executions to accomplish the execution of the overall order. Similarly, *Algorithmic Market Orders* refers to the situation when the client does not specify any price restrictions on the algorithm. Again, the algorithm then still uses combinations of individual market and limit orders to accomplish the execution of the overall order.

By using Algorithmic Limit Orders and thereby setting a lower or upper limit price threshold, the client explicitly allows that parts of the order do not have to be completed. It is obvious that execution performance of orders with such constraints should not be compared with the ACE Non-Discretionary estimates since it reflects some sort of opportunistic trading.

Because of the issues (i) and (ii) above, we consider two additional aggregation types:

1. aggregation by trade date, moniker id, stock id, side **and *Algorithmic Market Orders* only** (*Cluster Set 2*)
2. aggregation by trade date, moniker id, stock id, side **and *Algorithmic Market Orders* and structured strategies only, i.e., VWAP or IS** (*Cluster Set 3*)

The alternative Cluster Sets 2 and especially 3 can be considered as less biased with respect to points (i) and (ii) above since clients have not specified any execution limit price thresholds and VWAP and IS strategies are typically executing orders fully. From Table 1, it is apparent that the restrictions introduced for Cluster Sets 2 and 3 reduce the number of observations. While Cluster Set 1 has 1,119,460 observations, the number of observations in Cluster Set 2 is reduced to 810,920. However, Cluster Set 3 greatly reduces the number of observations to 92,890.

In what follows, we determine the average actual transaction cost for each of the three cluster types/sets and compare them to the associated ACE cost estimates. From the above discussion, one would assume that each of the three cluster sets has different trading instructions (and thus performance) and consequently should be compared to a different ACE benchmark (ACE Discretionary or ACE Non-Discretionary).

For each cluster of every cluster set, we determine the order start time, the average execution price, the total cluster size (relative to the corresponding daily MDV), the trading horizon (based on 30-minutes intervals), the prevailing mid quote prior to this time and the actual costs relative to this mid quote (in b.p.). In addition, we determine the average half-spread within the execution period as well as ACE Discretionary and ACE Non-Discretionary for the cluster. The ACE cost estimates are obtained using the actual trading strategy during the trading period. Next, we subdivide for each cluster set all clusters into Listed and Over-the-Counter (OTC) stocks to take into account cost differences for different market structures.[2] Stocks are then grouped based on their MDV. We rank all available stocks (approximately 7,000) according to their MDV at the beginning of each month during the sample period. For Listed and OTC stocks separately, we divide the stocks into three liquidity groups. The category "Least Liquid" represents the least liquid stocks, the category "Medium Liquid" represents stocks that are relatively liquid and the category "Most Liquid" consists of stocks that are traded very heavily.

Figures 1 to 6 below show the average actual costs by relative cluster size for the three cluster sets. As expected, average actual costs from the original aggregation (Cluster Set 1) are always the lowest. In addition, the difference in transaction costs is the largest for larger order sizes and/or least liquid stocks where opportunistic trading matters the most. As stocks become more liquid more opportunities arise even for orders that need to be completed in full and therefore bias problem (b) is less severe. Note that transaction costs are shown only up to 20% of MDV since our data sample has very few observations beyond that. Cluster Sets 2 and 3 appear to have very similar costs and there is no clear winner. This observation seems to be surprising since intuitively one would expect that the Cluster Set 2 should have lower transaction costs than Cluster Set 3 due to the inclusion of ITG Dark and Active Algorithms allowing for opportunistic trading. However, note that Figures 1 to 6 do not represent a proper comparison and we may have a bias in the selection of the clusters. The potential bias may be due to several things. First, the client selects the ITG Algorithm and thus may introduce a selection bias by matching the algorithm with order characteristics and market conditions. Second, independent of the client's decision, the characteristics of the underlying stocks of orders and the prevailing market conditions may be systematically different across different algorithms, again introducing a bias. To address these potential biases, we compare actual trading

[2] Listed stocks are listed on the New York Stock Exchange (NYSE) or the American Stock Exchange (Amex). All other stocks are considered OTC stocks.

costs with the different ACE cost estimates rather than simply across cluster sets in the remainder of the paper.

3. COMPARISON OF ACTUAL COSTS AND ACE DISCRETIONARY AND NON-DISCRETIONARY COST ESTIMATES

Figures 7 to 12 compare the actual average transaction costs for the Cluster Set 1 with the average ACE cost estimates (ACE Discretionary, ACE Non-Discretionary and ACE half-spread[3]) as well as with the spread cutoff a). The majority of the trades in Cluster Set 1 are from algorithms that are liquidity searching (Dark, Active, VP) and therefore (at least) partially opportunistic (Dark, Active).[4] In addition, a significant amount of clusters contain limit price thresholds (that is Algorithmic Limit Orders). As such it would be of no surprise if the average actual costs were very low, possibly lower than the average costs of a typical market participant that has a healthy mix between an opportunistic and deterministic trading style. Figures 7 to 12 support this intuition. Average ITG Algorithms costs are below ACE Discretionary that represents the average costs of all U.S. market participants. The opportunistic trading style of a substantial part of ITG Algorithms data set can be especially well observed in Figure 7. No matter how large the order size for illiquid stocks is, average ITG Algorithms costs are matching the half-spread. Note that the green spread threshold (cutoff a)) crosses the actual costs around an order size of 15% of MDV. Consequently, on average ITG Algorithms would be classified as opportunistic for orders above 15% of MDV for the less liquid stock universe. For Medium Liquid and Most Liquid stocks, actual trading costs and both ACE cost estimates are clearly higher than the average half- spread. It seems that the trading style is more balanced for these stock universes and order completion matters sometimes but not always. This is in line with the fact that the actual transaction costs of medium liquid and most liquid stocks are closer to ACE Discretionary than in the case of the Least Liquid stocks. One could argue that ITG Algorithms perform well measured against ACE Discretionary (which is the appropriate cost benchmark).

Overall, the above trading cost analysis for the entire ITG Algorithms data sample (Cluster Set 1) has shown that our intuition is correct. The ACE Discretionary benchmark is supposed to represent

[3] The ACE half-spread is the ACE cost estimate assuming price impact is zero.
[4] See Table 1 for the exact break-down.

typical costs of all market participants. Since we know that the majority of ITG Algorithms orders possess characteristics (i) and (ii), the actual ITG Algorithms costs should be lower than the benchmark. Our analysis confirms this. By eliminating Algorithmic Limit Orders (Cluster Set 2), we would expect to be closer to the ACE Discretionary framework. The intuition behind this is that by dropping Algorithmic Limit Orders, we eliminate orders with more discretion and as a result the mix of orders in our sample should exhibit costs closer to the typical costs of all market participants. We focus on this next.

Figures 13 to 18 compare the actual average transaction costs for Cluster Set 2 with the average ACE cost estimates (ACE Discretionary, ACE Non-Discretionary and ACE half-spread) as well as with the spread cutoff a). In contrast to Cluster Set 1, we have removed all orders for which the clients have selected limit price thresholds. As a consequence, opportunistic trading behavior is restricted by the internal usage of Dark, Active, and VP servers only and therefore one might assume that this sample is in line with the "average" existing trading behavior in the U.S. market. As such actual trading costs for Cluster Set 2 should be close to the ACE Discretionary benchmark.

Figures 13 to 18 confirm the intuition. The actual average ITG Algorithm costs for Cluster Set 2 are very close to the ACE Discretionary benchmark for most segments. In the few cases for which this is not the case (small order sizes and Medium Liquid, Listed stocks) the actual trading costs are still between the ACE Discretionary and ACE Non-Discretionary benchmark.

In what follows, we consider Cluster Set 3 which by construction is a good subsample to evaluate actual transaction costs for non-discretionary orders.

Figures 19 to 24 compare the actual average transaction costs for Cluster Set 3 with the average ACE cost estimates (ACE Discretionary, ACE Non-Discretionary and ACE half-spread) as well as with the spread cutoff a). As a reminder, the clusters in this data sample have been constructed using the structured strategies VWAP and IS, and Algorithmic Market Orders only. Because of these restrictions the data sample is limited and has shrunk considerably (we have less than 10% of the initial data, see Table 10). Therefore, actual trading costs are not as well-behaved as for the Cluster Sets 1 and 2.

ITG's structured strategies assume that each order is traded according to a fixed, pre-determined schedule. By definition, all orders will be completed fully. As a consequence, ACE Non-Discretionary should be an ideal benchmark for this existing subsample if our filtration of opportunistic clients in the peer group database is reasonable. In other words, we argue that ACE Non-Discretionary is a good benchmark for non-discretionary trading if actual trading costs for the Cluster Set 3 are close to the benchmark.

Figures 19 to 24 demonstrate exactly this behavior. Due to the limited amount of observations, actual trading costs have large fluctuations. Nevertheless, it appears that the actual costs fluctuate around the ACE Non-Discretionary benchmark for all scenarios. There is no bias observable. We conclude that our filtration heuristics for ACE Non-Discretionary works well and that the ACE Non-Discretionary benchmark is a meaningful benchmark for orders that need to be completed.

4. SUMMARY

We analyze the actual costs of ITG Algorithms data consisting of executions over the period from October 2006 to March 2007 and compare them with ACE cost estimates.

The following summarizes the results:

1. It appears that ITG Algorithms data overall is executed very opportunistically. This is particularly the case for larger order sizes and/or less liquid stocks. There is no suitable cost benchmark for opportunistic trading since any benchmark can be gamed.
2. By restricting ITG Algorithms data to Algorithmic Market Orders only, we see actual costs that are in line with typical average costs of market participants as reflected in the ACE Discretionary benchmark.
3. We test the ACE Non-Discretionary benchmark and look at actual transaction costs for orders that have been executed using structured strategies, i.e., strategies that are pre-determined before the trading begins. Orders that are executed through Implementation Shortfall (IS) or Volume-Weighted Average Price (VWAP) are typically filled 100%. The cost comparison of actual costs versus ACE Non-Discretionary shows that we are in the same ballpark. Thus, the heuristic filtration algorithm used for the calibration of ACE Non-Discretionary is reasonable and provides accurate estimates for all segments.

Table 1: ITG Algorithms Execution Data Sample

This table reports the number of observations for different ITG Algorithms in the ITG Algorithms Execution Data used in this study.

ITG Algorithm		All orders (Algorithmic Market and Limit Orders)	Algorithmic Market Orders only
Active	Number of observations	565,521	373,039
	% of total observations	50.52	46.00
Dark	Number of observations	353,256	264,758
	% of total observations	31.56	32.65
Volume Participation (VP)	Number of observations	97,005	80,233
	% of total observations	8.67	9.89
Volume-Weighted Average Price (VWAP)	Number of observations	91,075	82,480
	% of total observations	8.14	10.17
Implementation Shortfall (IS)	Number of observations	12,603	10,410
	% of total observations	1.13	1.28
Total		1,119,460	810,920

Figure 1: Actual Costs of Least Liquid, Listed Stocks

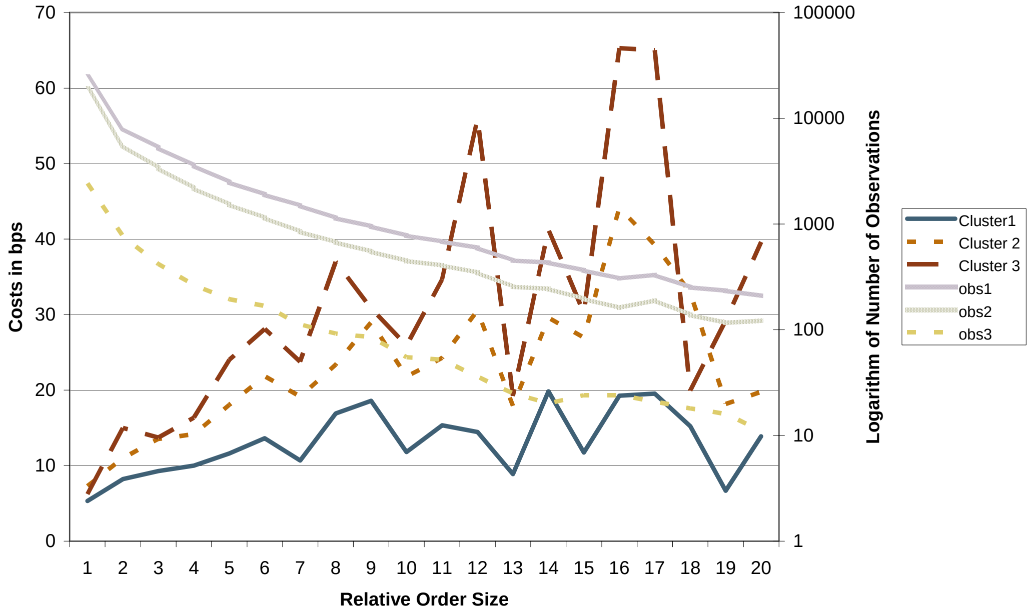

Cluster1, Cluster2, and Cluster3 denote Cluster Set 1, Cluster Set 2, and Cluster Set 3, respectively.
obs1, obs2, and obs3 denote number of observations for Cluster Set 1, Cluster Set 2, and Cluster Set 3, respectively.

Figure 2: Actual Costs of Medium Liquid, Listed Stocks

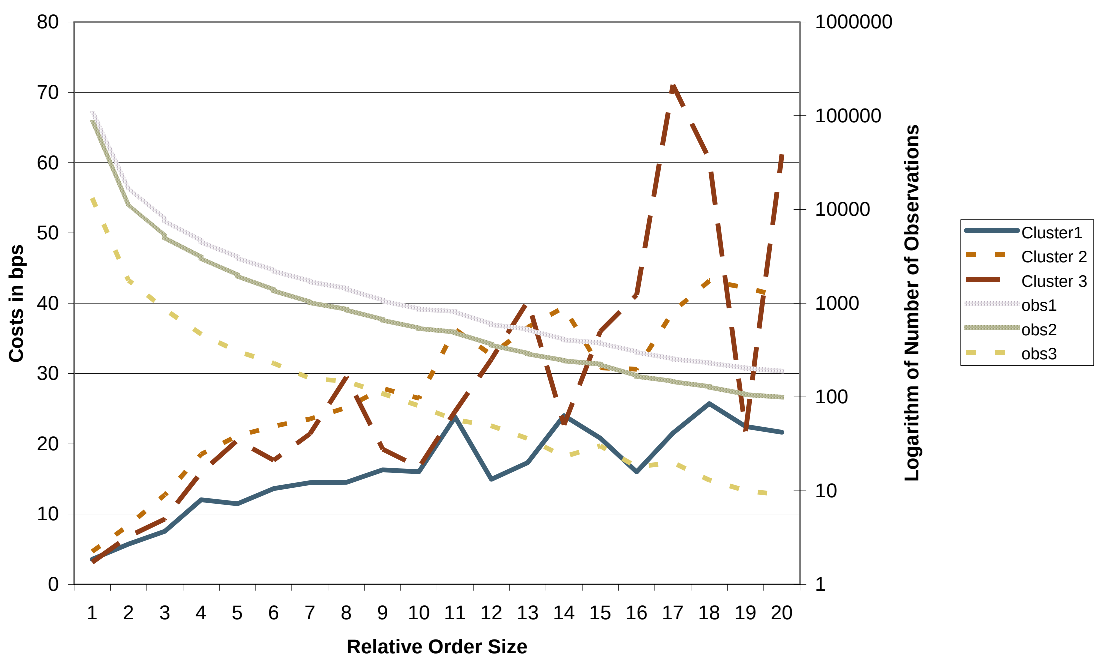

Cluster1, Cluster2, and Cluster3 denote Cluster Set 1, Cluster Set 2, and Cluster Set 3, respectively.
obs1, obs2, and obs3 denote number of observations for Cluster Set 1, Cluster Set 2, and Cluster Set 3, respectively.

Figure 3: Actual Costs of Most Liquid, Listed Stocks

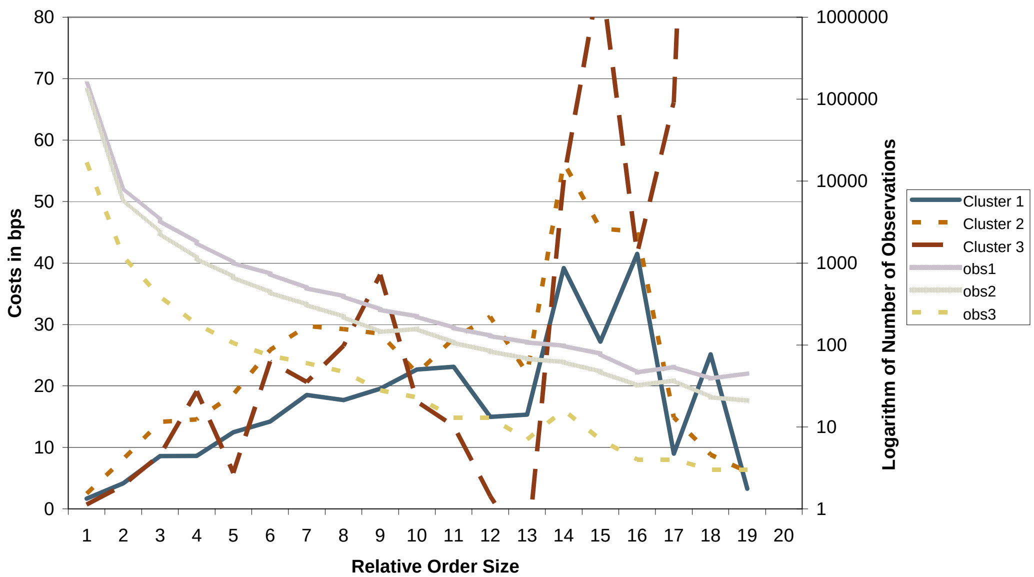

Cluster1, Cluster2, and Cluster3 denote Cluster Set 1, Cluster Set 2, and Cluster Set 3, respectively.
obs1, obs2, and obs3 denote number of observations for Cluster Set 1, Cluster Set 2, and Cluster Set 3, respectively.

Figure 4: Actual Costs of Least Liquid, OTC Stocks

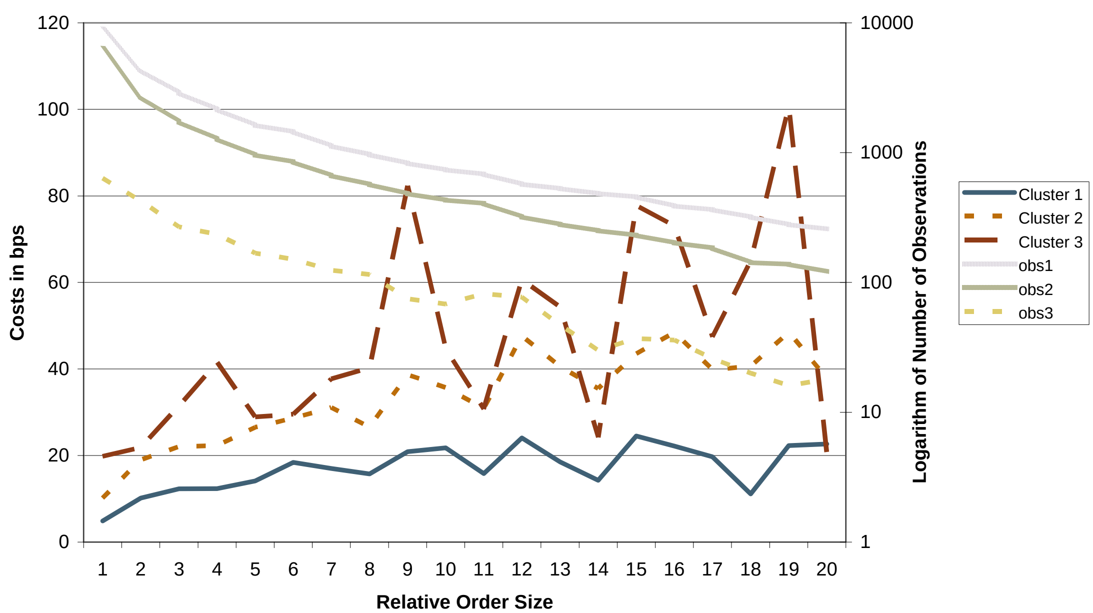

Cluster1, Cluster2, and Cluster3 denote Cluster Set 1, Cluster Set 2, and Cluster Set 3, respectively.
obs1, obs2, and obs3 denote number of observations for Cluster Set 1, Cluster Set 2, and Cluster Set 3, respectively.

Figure 5: Actual Costs of Medium Liquid, OTC Stocks

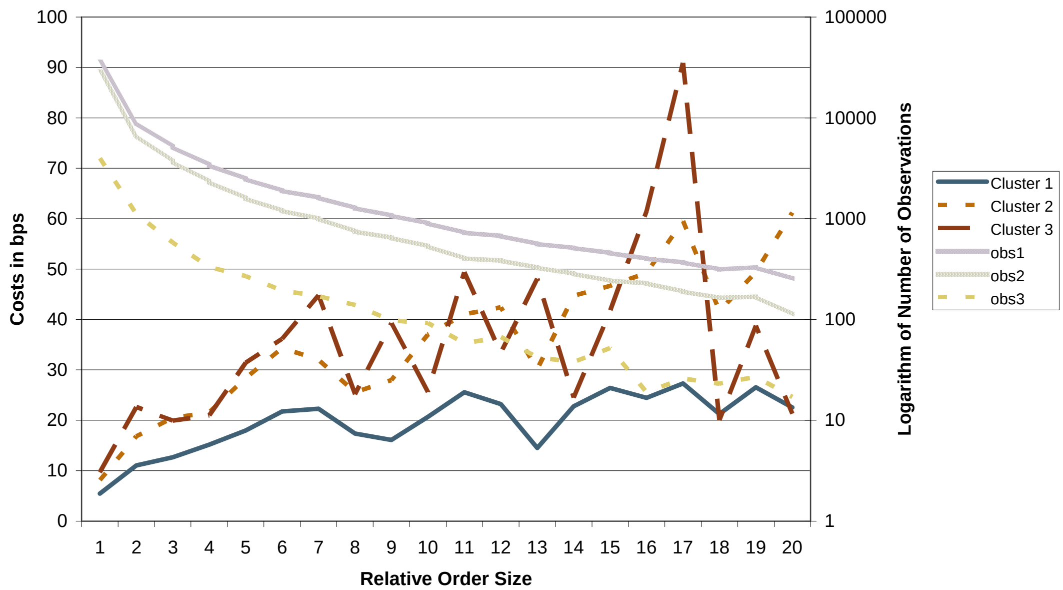

Cluster1, Cluster2, and Cluster3 denote Cluster Set 1, Cluster Set 2, and Cluster Set 3, respectively.
obs1, obs2, and obs3 denote number of observations for Cluster Set 1, Cluster Set 2, and Cluster Set 3, respectively.

Figure 6: Actual Costs of Most Liquid, OTC Stocks

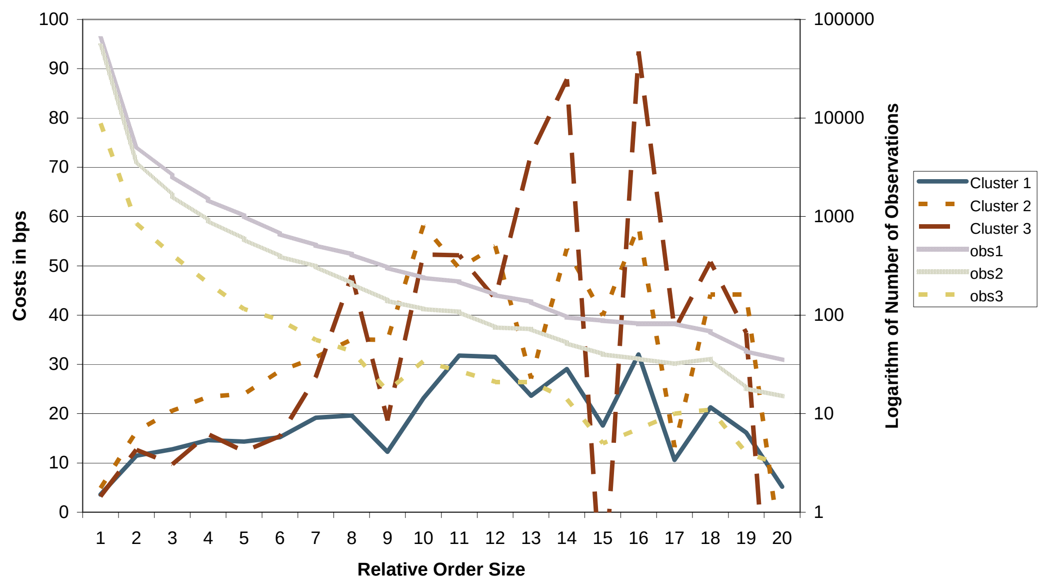

Cluster1, Cluster2, and Cluster3 denote Cluster Set 1, Cluster Set 2, and Cluster Set 3, respectively.
obs1, obs2, and obs3 denote number of observations for Cluster Set 1, Cluster Set 2, and Cluster Set 3, respectively.

Figure 7: Actual Costs Vs. ACE Costs for Less Liquid, Listed Stocks for Cluster Set 1 Using the Entire Data

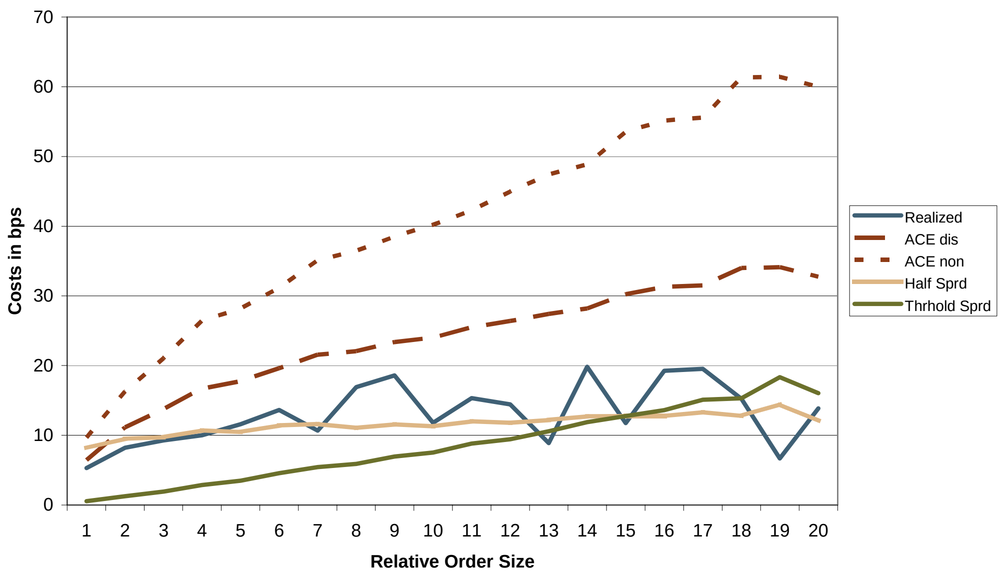

"Realized" denotes realized costs of the order. "ACE dis" denotes the ACE Discretionary cost estimates. "ACE non" denotes the ACE Non-Discretionary cost estimates. "Half Sprd" denotes the ACE cost estimates assuming a price impact of zero. "Thrhold Sprd" denotes the cutoff a) of the heuristics for the identification of opportunistic trading.

Figure 8: Actual Costs Vs. ACE Costs for Medium Liquid, Listed Stocks for Cluster Set 1 Using the Entire Data

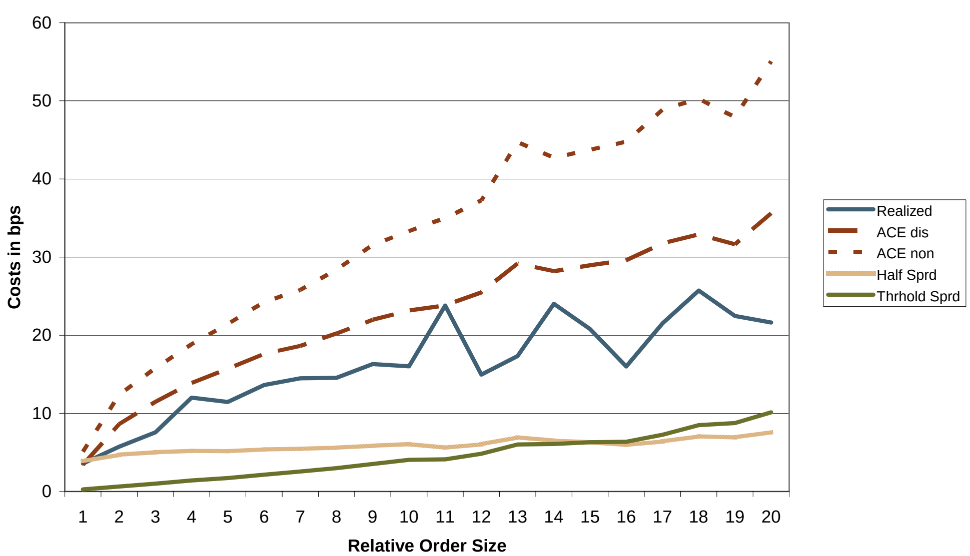

"Realized" denotes realized costs of the order. "ACE dis" denotes the ACE Discretionary cost estimates. "ACE non" denotes the ACE Non-Discretionary cost estimates. "Half Sprd" denotes the ACE cost estimates assuming a price impact of zero. "Thrhold Sprd" denotes the cutoff a) of the heuristics for the identification of opportunistic trading.

Figure 9: Actual Costs Vs. ACE Costs for Liquid, Listed Stocks for Cluster Set 1 Using the Entire Data

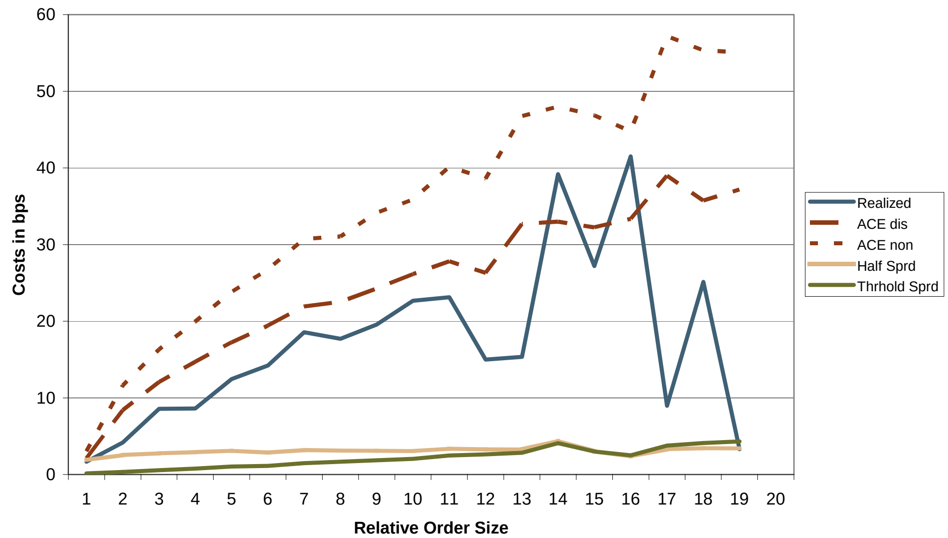

"Realized" denotes realized costs of the order. "ACE dis" denotes the ACE Discretionary cost estimates. "ACE non" denotes the ACE Non-Discretionary cost estimates. "Half Sprd" denotes the ACE cost estimates assuming a price impact of zero. "Thrhold Sprd" denotes the cutoff a) of the heuristics for the identification of opportunistic trading.

Figure 10: Actual Costs vs. ITG ACE Cost Estimates for Least Liquid, OTC Stocks for Cluster Set 1 Using the Entire Data

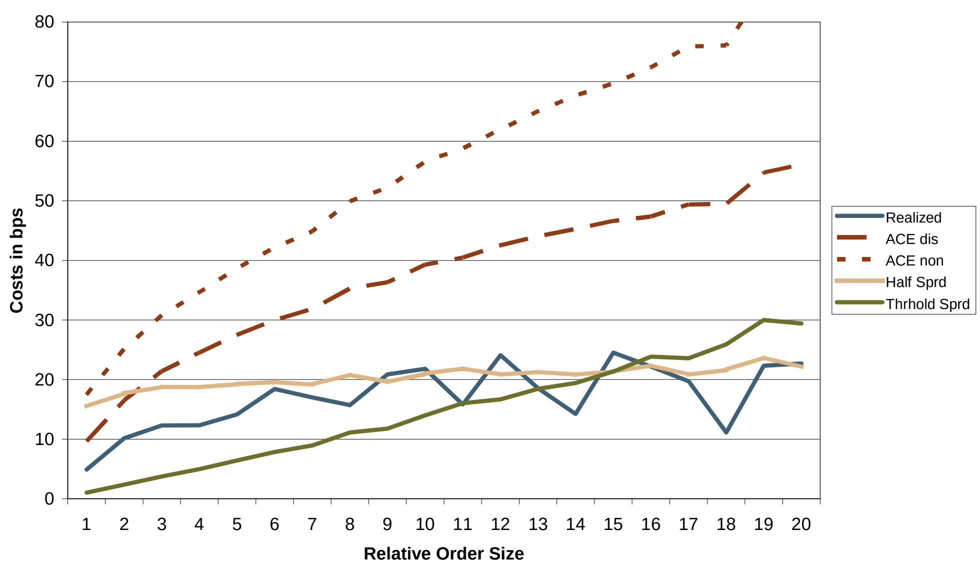

"Realized" denotes realized costs of the order. "ACE dis" denotes the ACE Discretionary cost estimates. "ACE non" denotes the ACE Non-Discretionary cost estimates. "Half Sprd" denotes the ACE cost estimates assuming a price impact of zero. "Thrhold Sprd" denotes the cutoff a) of the heuristics for the identification of opportunistic trading.

Figure 11: Actual Costs vs. ITG ACE Cost Estimates for Medium Liquid, OTC Stocks for Cluster Set 1 Using the Entire Data

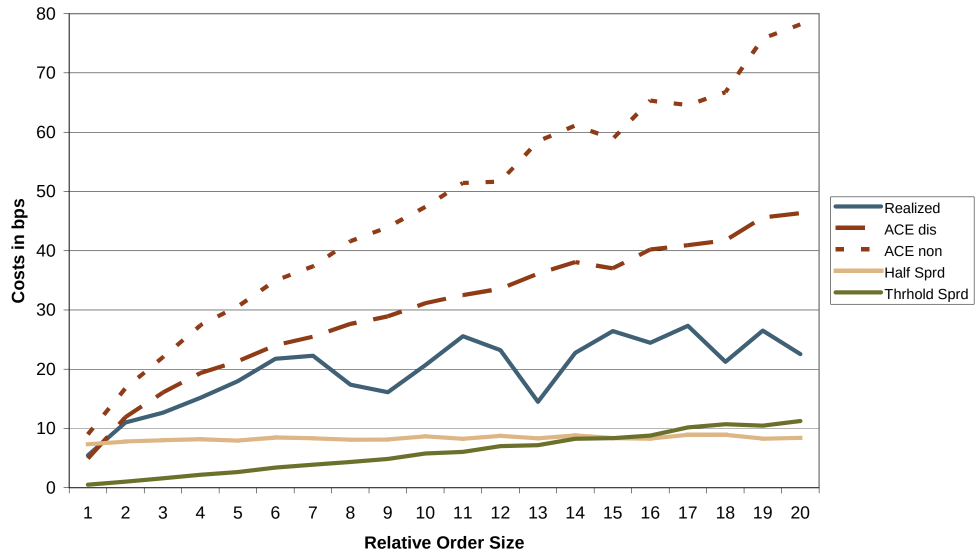

“Realized” denotes realized costs of the order. “ACE dis” denotes the ACE Discretionary cost estimates. “ACE non” denotes the ACE Non-Discretionary cost estimates. “Half Sprd” denotes the ACE cost estimates assuming a price impact of zero. “Thrhold Sprd” denotes the cutoff a) of the heuristics for the identification of opportunistic trading.

Figure 12: Actual Costs vs. ACE Cost Estimates for Most Liquid, OTC Stocks for Cluster Set 1 Using the Entire Data

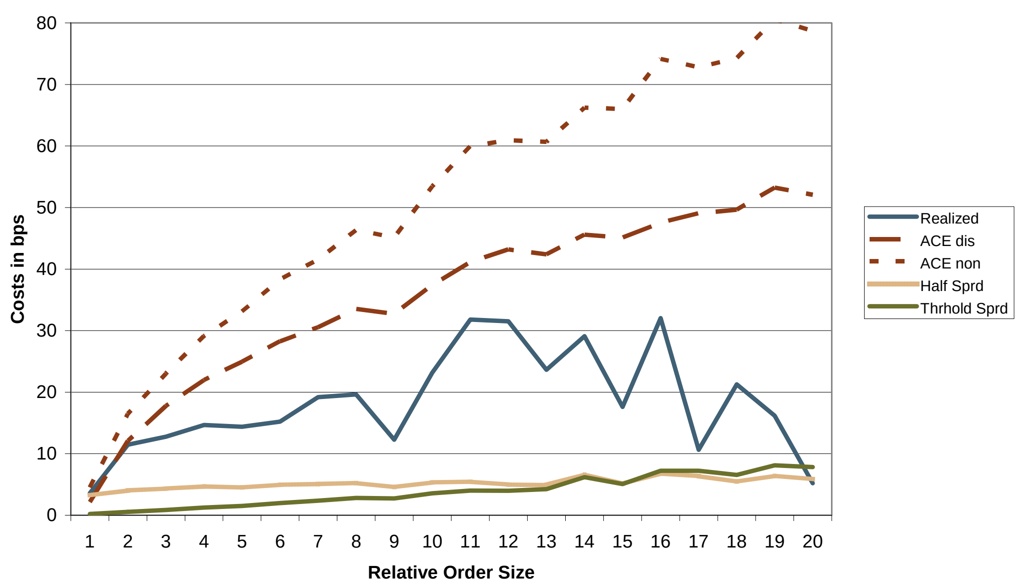

“Realized” denotes realized costs of the order. “ACE dis” denotes the ACE Discretionary cost estimates. “ACE non” denotes the ACE Non-Discretionary cost estimates. “Half Sprd” denotes the ACE cost estimates assuming a price impact of zero. “Thrhold Sprd” denotes the cutoff a) of the heuristics for the identification of opportunistic trading.

Figure 13: Actual Costs vs. ACE Cost Estimates for Least Liquid, Listed Stocks for Cluster Set 2

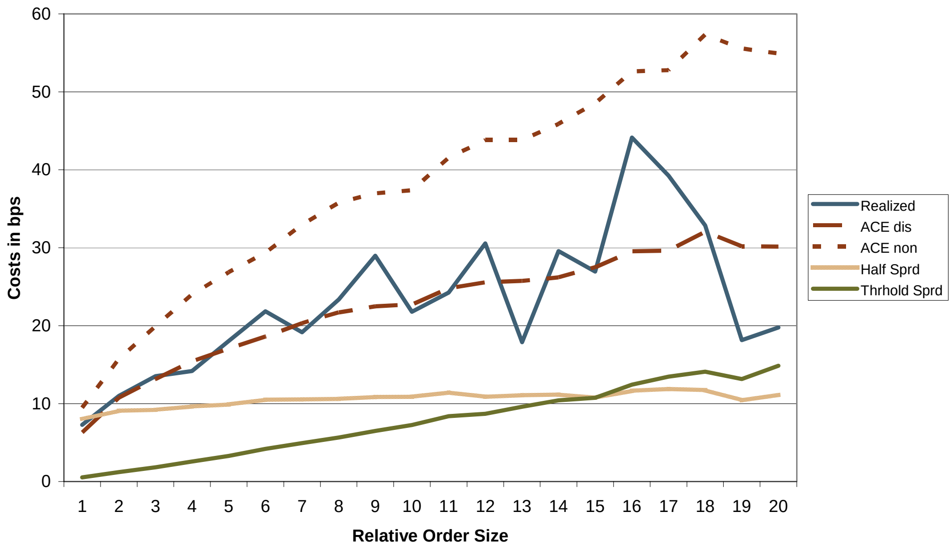

"Realized" denotes realized costs of the order. "ACE dis" denotes the ACE Discretionary cost estimates. "ACE non" denotes the ACE Non-Discretionary cost estimates. "Half Sprd" denotes the ACE cost estimates assuming a price impact of zero. "Thrhold Sprd" denotes the cutoff a) of the heuristics for the identification of opportunistic trading.

Figure 14: Actual Costs vs. ACE Cost Estimates for Medium Liquid, Listed Stocks for Cluster Set 2

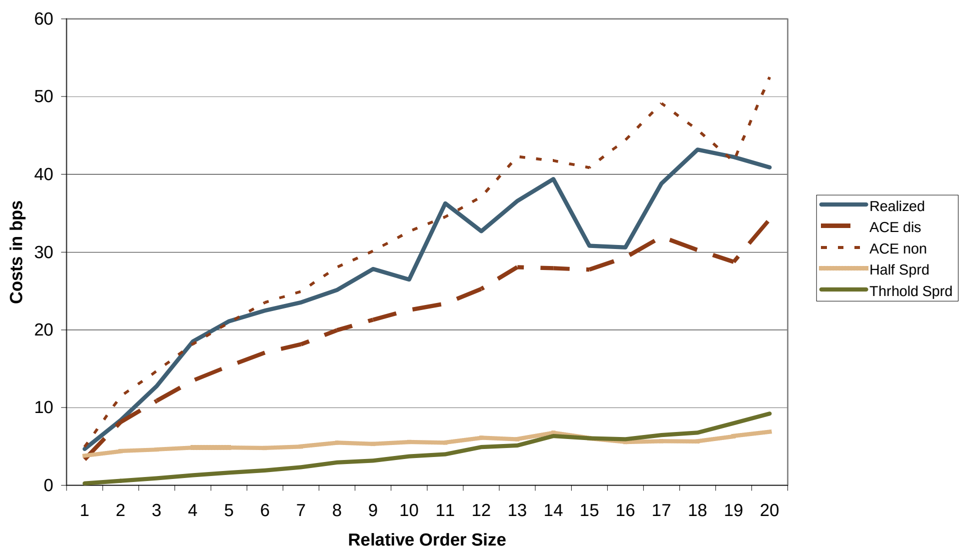

"Realized" denotes realized costs of the order. "ACE dis" denotes the ACE Discretionary cost estimates. "ACE non" denotes the ACE Non-Discretionary cost estimates. "Half Sprd" denotes the ACE cost estimates assuming a price impact of zero. "Thrhold Sprd" denotes the cutoff a) of the heuristics for the identification of opportunistic trading.

Figure 15: Actual Costs vs. ACE Cost Estimates for Most Liquid, Listed Stocks for Cluster Set 2

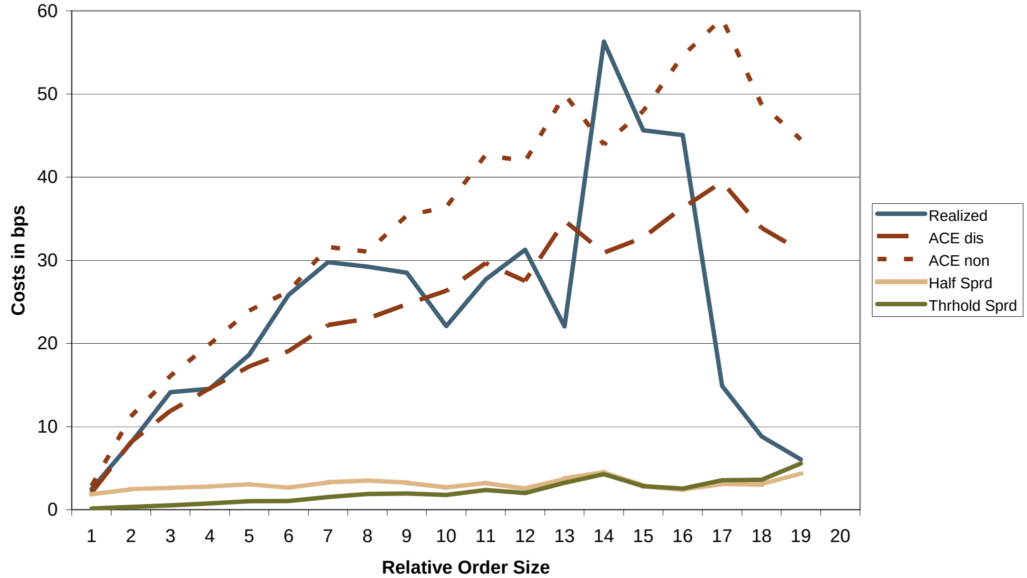

“Realized” denotes realized costs of the order. “ACE dis” denotes the ACE Discretionary cost estimates. “ACE non” denotes the ACE Non-Discretionary cost estimates. “Half Sprd” denotes the ACE cost estimates assuming a price impact of zero. “Thrhold Sprd” denotes the cutoff a) of the heuristics for the identification of opportunistic trading.

Figure 16: Actual Costs vs. ACE Cost Estimates for Least Liquid, OTC Stocks for Cluster Set 2

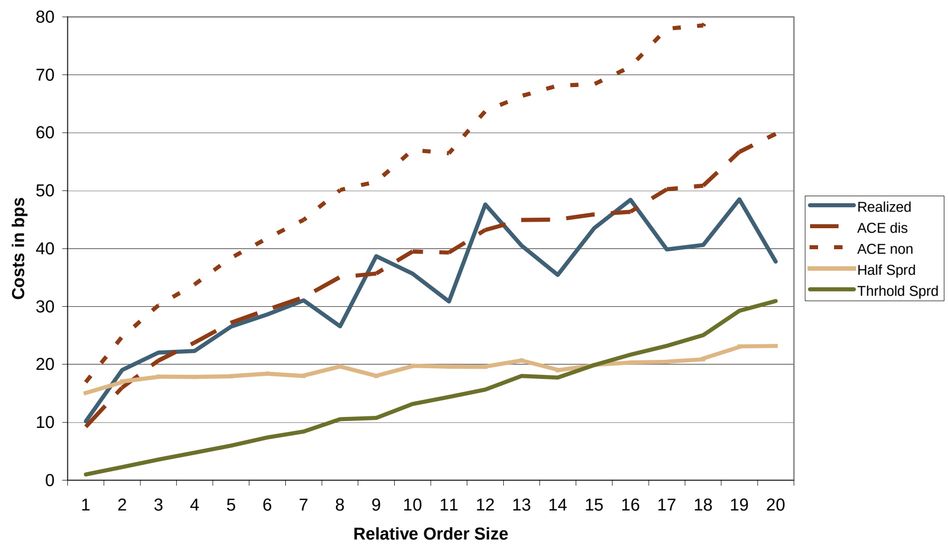

“Realized” denotes realized costs of the order. “ACE dis” denotes the ACE Discretionary cost estimates. “ACE non” denotes the ACE Non-Discretionary cost estimates. “Half Sprd” denotes the ACE cost estimates assuming a price impact of zero. “Thrhold Sprd” denotes the cutoff a) of the heuristics for the identification of opportunistic trading.

Figure 17: Actual Costs vs. ACE Cost Estimates for Medium Liquid, OTC Stocks for Cluster Set 2

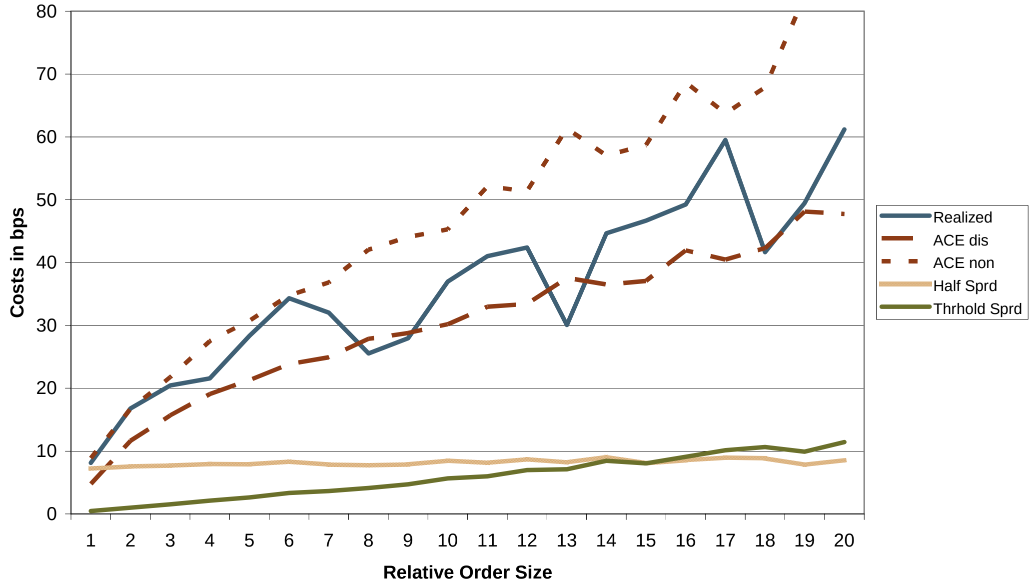

“Realized” denotes realized costs of the order. “ACE dis” denotes the ACE Discretionary cost estimates. “ACE non” denotes the ACE Non-Discretionary cost estimates. “Half Sprd” denotes the ACE cost estimates assuming a price impact of zero. “Thrhold Sprd” denotes the cutoff a) of the heuristics for the identification of opportunistic trading.

Figure 18: Actual Costs vs. ACE Cost Estimates for Most Liquid, OTC Stocks for Cluster Set 2

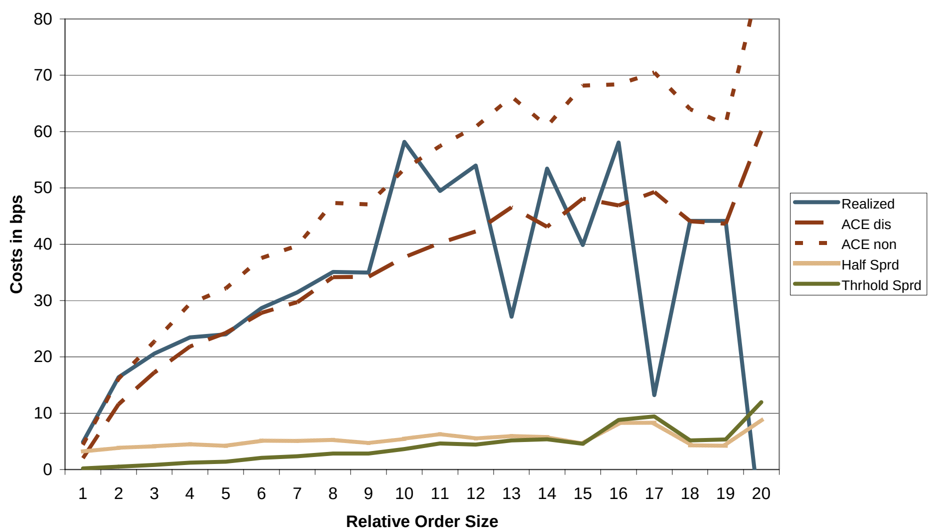

“Realized” denotes realized costs of the order. “ACE dis” denotes the ACE Discretionary cost estimates. “ACE non” denotes the ACE Non-Discretionary cost estimates. “Half Sprd” denotes the ACE cost estimates assuming a price impact of zero. “Thrhold Sprd” denotes the cutoff a) of the heuristics for the identification of opportunistic trading.

Figure 19: Actual Costs vs. ACE Cost Estimates for Least Liquid, Listed Stocks for Cluster Set 3

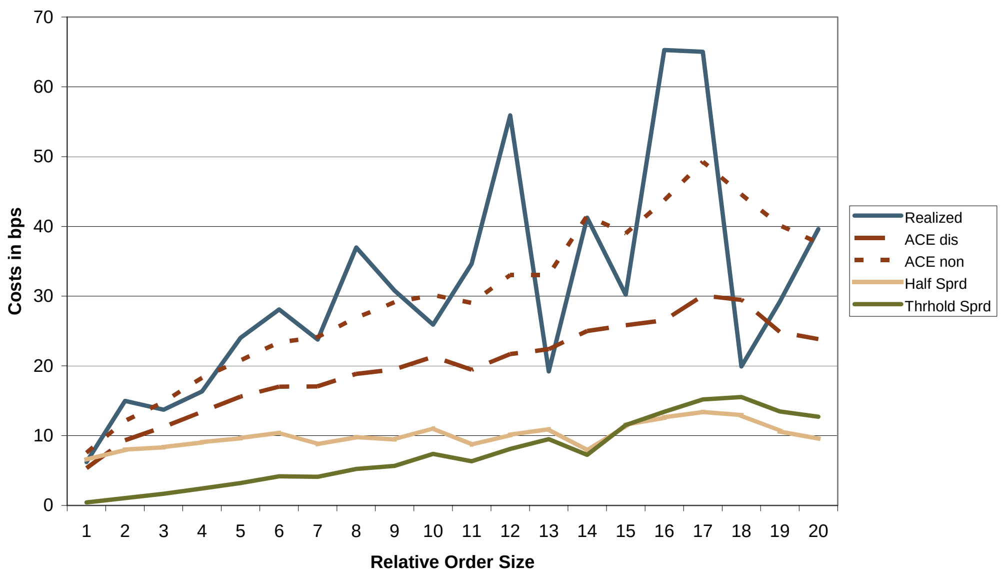

“Realized” denotes realized costs of the order. “ACE dis” denotes the ACE Discretionary cost estimates. “ACE non” denotes the ACE Non-Discretionary cost estimates. “Half Sprd” denotes the ACE cost estimates assuming a price impact of zero. “Thrhold Sprd” denotes the cutoff a) of the heuristics for the identification of opportunistic trading.

Figure 20: Actual Costs vs. ACE Cost Estimates for Medium Liquid, Listed Stocks for Cluster Set 3

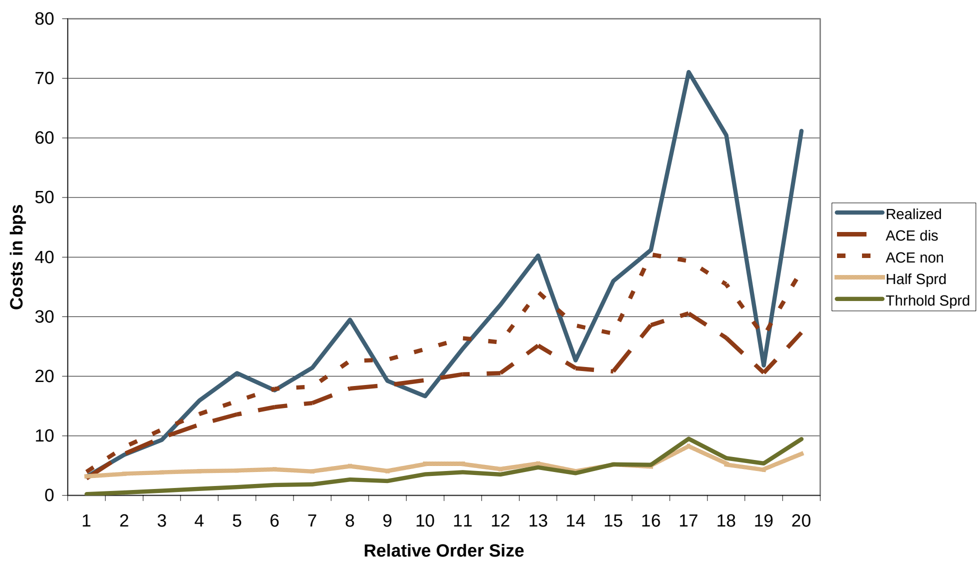

“Realized” denotes realized costs of the order. “ACE dis” denotes the ACE Discretionary cost estimates. “ACE non” denotes the ACE Non-Discretionary cost estimates. “Half Sprd” denotes the ACE cost estimates assuming a price impact of zero. “Thrhold Sprd” denotes the cutoff a) of the heuristics for the identification of opportunistic trading.

Figure 21: Actual Costs vs. ACE Cost Estimates for Most Liquid, Listed Stocks for Cluster Set 3

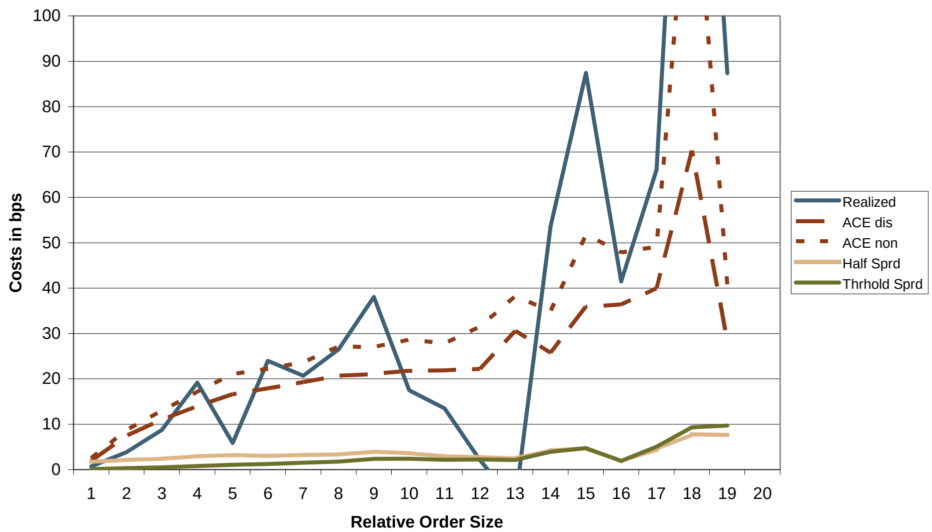

"Realized" denotes realized costs of the order. "ACE dis" denotes the ACE Discretionary cost estimates. "ACE non" denotes the ACE Non-Discretionary cost estimates. "Half Sprd" denotes the ACE cost estimates assuming a price impact of zero. "Thrhold Sprd" denotes the cutoff a) of the heuristics for the identification of opportunistic trading.

Figure 22: Actual Costs vs. ACE Cost Estimates for Least Liquid, OTC Stocks for Cluster Set 3

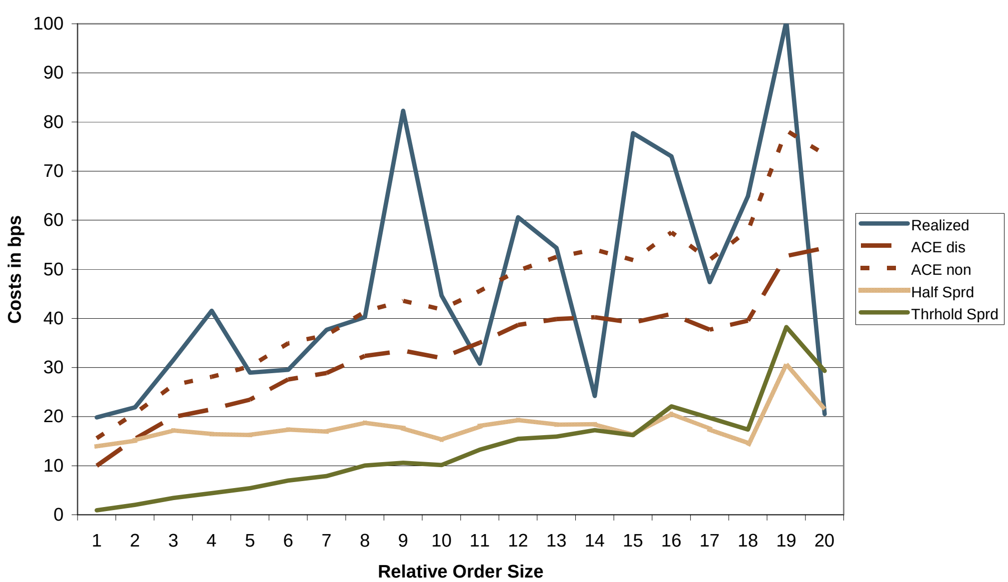

"Realized" denotes realized costs of the order. "ACE dis" denotes the ACE Discretionary cost estimates. "ACE non" denotes the ACE Non-Discretionary cost estimates. "Half Sprd" denotes the ACE cost estimates assuming a price impact of zero. "Thrhold Sprd" denotes the cutoff a) of the heuristics for the identification of opportunistic trading.

Figure 23: Actual Costs vs. ACE Cost Estimates for Medium Liquid, OTC Stocks for Cluster Set 3

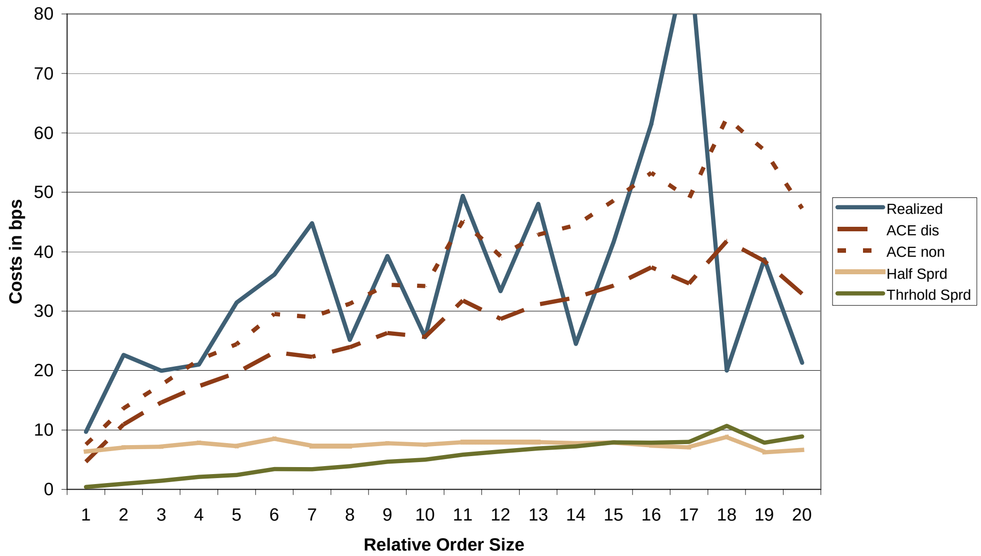

“Realized” denotes realized costs of the order. “ACE dis” denotes the ACE Discretionary cost estimates. “ACE non” denotes the ACE Non-Discretionary cost estimates. “Half Sprd” denotes the ACE cost estimates assuming a price impact of zero. “Thrhold Sprd” denotes the cutoff a) of the heuristics for the identification of opportunistic trading.

Figure 24: Actual Costs vs. ACE Cost Estimates for Most Liquid, OTC Stocks for Cluster Set 3

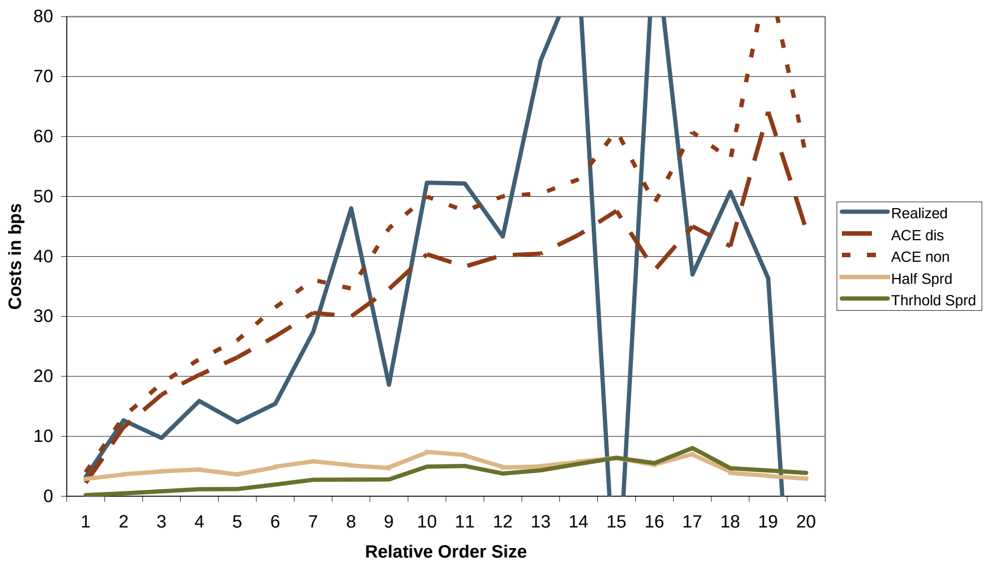

“Realized” denotes realized costs of the order. “ACE dis” denotes the ACE Discretionary cost estimates. “ACE non” denotes the ACE Non-Discretionary cost estimates. “Half Sprd” denotes the ACE cost estimates assuming a price impact of zero. “Thrhold Sprd” denotes the cutoff a) of the heuristics for the identification of opportunistic trading.