

Measuring and Modeling Execution Cost and Risk¹

Robert Engle
NYU and Morgan Stanley

Robert Ferstenberg
Morgan Stanley

Jeffrey Russell
University of Chicago

April 2006

Preliminary
Please do not quote without authors' permission

Abstract:

We introduce a new analysis of transaction costs that explicitly recognizes the importance of the timing of execution in assessing transaction costs. Time induces a risk/cost tradeoff. The price of immediacy results in higher costs for quickly executed orders while more gradual trading results in higher risk since the value of the asset can vary more over longer periods of time. We use a novel data set that allows a sequence of transactions to be associated with individual orders and measure and model the expected cost and risk associated with different order execution approaches. The model yields a risk/cost tradeoff that depends upon the state of the market and characteristics of the order. We show how to assess liquidation risk using the notion of liquidation value at risk (LVAR).

¹ This paper is the private opinion of the authors and does not necessarily reflect policy or research of Morgan Stanley. We thank Peter Bolland for very useful discussions as well as seminar participants at the University of Pennsylvania and the NYSE Economics Research Group. Jeffrey Russell gratefully acknowledges Morgan Stanley and NYU for funding a visiting position at NYU Stern School of Business where this research was conducted.

1. Introduction

Understanding execution costs has important implications for both practitioners and regulators and has attracted substantial attention from the academic literature. Traditional analysis of transaction costs focus on the average distance between observed transaction prices and an “efficient” or fair market price. These types of analysis, however, are disconnected from transaction costs faced in practice since they neglect any notion of risk. Specifically, a buy order could be filled by submitting a market order and paying a price near the ask. Alternatively, the order could be submitted as a limit order and either execute at a better price, or not execute at all. Similarly, a single order is often broken up into a sequence of smaller ones spread out over time. This temporal dimension to the problem yields a natural cost/risk tradeoff. Orders executed over a short period of time will have a high expected cost associated with immediate execution but the risk will be low since the price is (nearly) known immediately. Orders executed over a long period of time may have a smaller price impact and therefore smaller expected cost but may be more risky since the asset price can vary more over longer periods of time than shorter periods of time. Using a novel data set that allows transactions to be associated with individual orders we measure and model the expected cost and risk associated with different order execution strategies.

Our empirical work builds directly on the recent research of Almgren and Chriss (1999, 2000), Almgren (2003), Grinold and Kahn (1999), Obizhaeva and Wang (2005), and Engle and Ferstenberg (2006). These papers examine execution quality involving not just the expected cost but also the risk dimension. Order execution strategies that are guaranteed to execute quickly offer a different risk/reward tradeoff than transaction

strategies that can take a longer time to be filled. The result is a frontier of risk/reward tradeoffs that is familiar in finance and analogous to classic mean variance analysis of portfolios. In fact, the work of Engle and Ferstenberg (2006) show that this analogy is deeper than might appear at first glance. Namely, they show how to integrate the portfolio decision and execution decision into a single problem and how to optimize these choices jointly.

Our work differs in important ways from most traditional approaches to the analysis of transaction costs. The classic measures of transaction costs such as Roll's measure, (realized) effective spreads, or the half spread measure average (positive) deviations of transaction prices from a notional efficient price². The midquote is often taken as the efficient price. As such, these measures focus purely on expected cost and are not well suited to analyze the cost of limit order strategies or the splitting up of orders into smaller components. Part of the limitations of the traditional analysis of transaction costs is driven by data availability. Standard available data does not generally include information about how long it took before a limit order executed. Even more rare is information providing a link between individual trades and the larger orders.

Using a unique data set consisting of 233,913 orders executed by Morgan Stanley in 2004, we are able to construct measures of both the execution risk and cost³. Our data includes information about when the order was submitted and the times, prices, and quantities traded in filling the order. This data allows to take a novel view of the costs and risks associated with order execution.

The expected cost and variance tradeoffs that the trader faces will depend upon the liquidity conditions in the market and the characteristics regarding the order. We model both the expected cost and the risk as a function of a series of conditioning variables. In this way, we are able to generate a time varying menu of expected cost and risk tradeoffs given the state of the market and order characteristics. The result is a conditional frontier

² For a survey of the literature see the special issue on transaction costs in the Journal of Financial Markets.

³ We do not know the identities of the traders and the data never left the confines of Morgan Stanley.

of different cost/risk tradeoffs. This frontier represents a menu of expected cost and variance tradeoffs faced by the trader.

The paper is organized as follows. Section 2 discusses measuring the order execution cost and risk. Section 3 presents the data used in our analysis and some preliminary analysis. Section 4 presents a model for conditional cost and risk with estimates. Section 5 presents an application of the model to liquidity risk and finally, section 6 concludes.

2. Measuring order execution cost and risk.

Our measure of trading costs captures both the expected cost and risk of execution. A key element of the measure takes the price available at the time of order submission as the benchmark price. The order may be executed using a larger number of small trades. Each transaction price and quantity traded might be different. The cost of the trade is always measured relative to a benchmark price which is taken to be the price available at the time of order submission. The transaction cost measure is then a weighted sum of the difference between the transaction price and the benchmark arrival price where the weights are simply the quantities traded. See Chan and Lakonishok (1995), Grinold and Kahn (1999), Almgren and Chriss (1999, 2000), Bertismas and Lo (1998) among others. In this paper the term order refers to the total volume that the agent desires to transact. We will use the term transaction to refer to a single trade. An order may be filled using multiple transactions.

More formally, let the position measured in shares at the end of time period t be x_t so that the number of shares transacted in period t is simply the change in x_t . Let p_0 denote the fair market value of the asset at the time of the order arrival. This can be taken to be the midquote at the time of the order arrival for this price in practice. Let \tilde{p}_t denote the transaction price of the asset in period t . The transaction cost for a given order is then given by:

$$(1) \quad TC = \sum_{t=1}^T \Delta x_t (\tilde{p}_t - p_0)$$

If the order is purchasing shares then the change in the number of shares will be non-negative. Transactions that occur above the reference price will therefore contribute positively toward transaction costs. Alternatively, when liquidating shares, the change in shares will be non-positive. Transaction prices that occur below the reference price will therefore contribute positively to transaction costs. For a given order, the transaction costs can be either negative or positive depending upon whether the price moved with or against the direction of the order. However, because each trade has a price impact that tends to move the price up for buys and down for sells we would expect the transaction cost to be positive on average. Given transaction cost, both a mean and variance of the transaction cost can be constructed.

Of course, a measure of the transaction cost per dollar traded is obtained by dividing the transaction cost by the arrival value:

$$(2) \quad TC\% = \frac{TC}{(x_T - x_0)P_0}$$

This measure allows for more meaningful comparison of costs across different orders and it used in our analysis.

The transaction cost can be decomposed into two components that provide some insight. Specifically, the transaction cost can be written as

$$(3) \quad TC = \sum_{t=1}^T \Delta x_t (\tilde{p}_t - p_t) + \sum_{t=1}^T (x_t - x_{t-1}) \Delta p_t$$

The first term represents the deviation of the transaction price from the local arrival price p_t . The former is closely related to traditional measures of transaction costs capturing local effects. The second term captures an additional cost due to the price impact. Each trade has the potential to move the value of the asset. This change in the asset price has an effect on all subsequent trades executed. Since the price impact typically moves the price to a less desirable price for the trader, this term will generally increase the cost of executing an order that would be missed by traditional measures that lack this temporal component.

3. The data

In order to analyze our transaction cost measure we need detailed order execution data that includes the arrival price, trade sizes and transaction prices that associated with all the transactions that were used to fill a given order. We obtained such data from Morgan Stanley. We do not know the identity of the traders that placed these orders and more importantly we do not know their motives. The orders could have been initiated by Morgan Stanley traders on behalf of their clients or by a buy side trader on behalf of a portfolio manager. Regardless, we do not know the identity of the Morgan Stanley trader or the client. We use the word “trader” to refer to either one. The order data never left the confines of Morgan Stanley and will not be made available outside of the confines of Morgan Stanley.

The orders were executed by Morgan Stanley’s Benchmark Execution Strategies™ (BXS) strategies during 2004. BXS is a order execution strategy that minimizes the expected cost of the trade for a given level of risk relative to a benchmark. The trades are “optimally” chosen relying on an automated trading procedure that specifies when and how much to trade. The algorithm changes the trading trajectory as the current trading conditions in the market vary⁴.

We consider two types of orders. The arrival price (AP) strategy and the volume weighted average price (VWAP). The AP strategy attempts to minimize the cost for a given level of risk around the arrival price p_0 . The trader can specify a level of urgency given by high, medium, and low urgency. The level of urgency is inversely related to the level of risk that the trader is willing to tolerate. High urgency orders have relatively low risk, but execute at a higher average cost. The medium and low urgency trades execute with progressively higher risk but at a lower average cost. The trader chooses the urgency and the algorithm derives the time to complete the trade given the state of the market and the trader’s constraints. For a given order size and market conditions, lower

⁴ The trading algorithm is a variant of Almgren and Chriss (2000) and the interested reader is referred to this paper for more details.

urgency orders tend to take longer to complete than higher urgency trades. However, since the duration to completion depends upon the market conditions and other factors there is not perfect correspondence between the urgency level and the time to complete the order. All orders in our sample, regardless of urgency, are filled within a single day.

We also consider VWAP orders. For these orders, the trader selects a time horizon and the algorithm attempts to execute the entire order by trading proportional to the market volume over this time interval. We only consider VWAP orders where the trader directed that the order be filled over the course of the entire trading day or that the overall volume traded was a very small fraction of the market volume over that period. This can be interpreted as a strategy to minimize cost regardless of risk. As such, we consider this a risk neutral trading VWAP strategy. Generally, these orders take longer to fill than the low urgency orders and should provide the highest risk and the lowest cost.

We consider orders for both NYSE and NASDAQ stocks. In order to ensure that orders of a given urgency reflect the cost/risk tradeoff optimized by the algorithm we apply several filters to the orders. Only completed orders are considered. Hence orders that begin to execute and are then cancelled midstream are not included in order to ensure homogeneity of orders of a given type. We excluded short sales because the uptick rule prevents the economic model from being used "freely". We do not consider orders executed prior to 9:36 since the market conditions surrounding the open are quite different than non-opening conditions. Only stocks that have an arrival price greater than \$5 are included. Orders that execute in less than 5 minutes tend to be very small orders that may be traded in a single trade. As such, they are not representative of the cost/risk tradeoff optimized by the algorithm. For similar reasons, orders smaller than 1000 shares are also not included. Finally, orders that are constrained to execute more quickly than the algorithm would dictate due to the approaching end of the trading day are also excluded. In the end, we are left with 233,913 orders.

For each order we construct the following statistics. The percent transaction cost are constructed using equation (2). The 5 day lagged bid ask spread weighted by time as a

percent of the midquote. The annualized 21 day lagged close to close volatility. The order shares divided by the lagged 21 day median daily volume. Table 1 presents summary statistics of our data. The statistics weight each order by its fraction of dollar volume. The transaction cost standard deviation is large relative to the average cost. Hence the risk component appears to be substantial.

The rows labeled B and S break down the orders into buyer and seller orders respectively. 62% of the dollars traded were buys and 38% sells. Buy orders tend to be slightly more expensive on average in this sample. The risk is similar. We see that 75% of the dollar volume was for NYSE stocks and 25% for Nasdaq. We see that NYSE orders tend to cost less than NASDAQ by an average of about 5 basis points. It is important to note that these statistics are unconditional and do not control for differences in characteristics of the stocks traded on the two exchanges which might be driving some of the variation in the observed costs. For example, we see that the average volatility of NASDAQ stocks is substantially higher than that of NYSE.

The last four rows separate the orders by urgency. H, M, and L, correspond to high, medium and low urgencies and V is the VWAP strategy. Hence as we move down the rows we move from high cost, low risk strategies to low cost, high risk strategies. Almost half of the orders are the risk neutral VWAP strategy (46%). Only 10% of the order volume is high urgency, 24% is medium urgency and 20% is high urgency. This is reflected in the sample statistics. The average cost decreases from 11.69 basis points to 8.99 basis points as we move from high to low urgency orders. At the same time, the risk moves from 12.19 basis points up to 40.89 basis points for the same change in urgency. Contrary to the intent, the VWAP strategy does not exhibit the lowest cost at 9.69 basis points. It is the most risky however. Of course, the order submission may depend on the state of the market and characteristics of the order. These unconditional statistics will not reflect the market state and may blur the tradeoffs faced by the traders.

Table 2 presents the same summary statistics conditional on the size of the order relative to the 21 day median daily volume. The first bin is for orders less than a quarter of a

percent of the 21 day median daily volume and the largest bin considered is for orders that exceed 1%. For each bin the statistics are presented for each type of order. The top of the table is for NYSE and the bottom half of the table is for NASDAQ stocks. Not surprisingly, for each order type, larger order sizes tend to be associated with higher average cost and higher risk indicating that larger orders are more difficult to execute along both the cost and risk dimensions.

This tradeoff can be seen clearly by plotting the cost/risk tradeoffs for each of the percent order size bins. Figure 1 presents the average cost/risk tradeoff for the NYSE stocks. Each contour indicates the expected cost/risk tradeoff faced for a given order size. Each contour is constructed using 4 points, the three urgencies and the VWAP. For a given contour, as we move from left to right we move from the high urgency orders to the VWAP. Generally speaking, the expected cost falls as the risk increases. This is not true for every contour, however. Increasing the percent order size shifts the entire frontier toward the north east indicating a less favorable average cost / risk tradeoff. Figure 2 presents the same plot but for NASDAQ stocks.

Contrary to what might be expected, some order size bins exhibit a cost increase as we move to less urgent strategies. These plots, however, do not consider the state of the market at the time the order is executed. It is entirely possible that the traders consider the state of the market when considering what type of urgency to associate with their order. If this is the case, a more accurate picture of the tradeoff faced by the trader can be obtained by considering the conditional frontier. This requires building a model for the expected cost and the standard deviation of the cost conditional on the state of the market. This is precisely the task considered in the next section of the paper.

4. Modeling the expected cost and risk of order execution.

Both the transaction cost and risk associated with trading a given order will vary depending on the state of the market. In this section we propose a modeling strategy for

both the expected cost and risk of trading an order. The model is estimated using the Morgan Stanley execution data described in the previous section. In estimating this model for the expected cost and risk we are also estimating a conditional expected cost/risk frontier. This frontier depicts the expected cost/risk tradeoff faced by the agent given the current state of the market. This frontier will be a function of both the state of the market as well as the size of the order.

Both the mean and the variance of transaction costs are assumed to be an exponential function of the market variables and the order size. Specifically the transaction costs for the i^{th} order are given by:

$$(4) \quad TC\%_i = \exp(X_i\beta) + \exp\left(\frac{1}{2} X_i\gamma\right) \varepsilon_i$$

where $\varepsilon_i \sim iid N(0,1)$. The conditional mean is an exponential function of a linear combination of the X_i with parameter vector β . The conditional standard deviation is also an exponential function of a linear combination of the X_i with parameter vector γ . X_i is a vector of conditioning information. In our empirical work we find that the same vector X_i explains both the mean and the variance but this restriction is obviously not required.

The exponential specification for both the mean and the variance restricts both to be positive numbers. This is a natural restriction for both the mean and the variance. While the realized transaction cost for any given trade can be either positive or negative (and empirically we do find both signs), the *expected* transaction cost is positive.

We consider several factors that market microstructure theory predicts should contribute to the ease of executing a given order. The lagged 5 day time weighted average spread as a percent of the midquote. The log volatility constructed from the average close to close returns over the last 21 days. The log of the average historical 21 day median daily dollar volumes. In addition to these market variables we also condition on the log of the dollar value of the order and the urgency associated with the order. The urgency is captured by 3 dummy variables for high, medium, and low urgencies. The constant term in the mean and variance models therefore corresponds to the VWAP strategy.

The exponential specification for the mean is not commonly used in econometrics analysis. It is particularly useful here since it is natural to restrict the mean to be positive. The often used method of modeling the logarithm of the left hand side variable won't work here because the transaction cost often take negative values. Also, notice that $\ln[E(TC\%_i)] = X_i\beta$. Hence the coefficients can be interpreted as the percent change in TC% for a one unit change in X. Right hand side variables that are expressed as the logarithm of a variable (such as $\ln(\text{value})$) can be interpreted as an elasticity with respect to the non-logged variable (such as value).

The exponential model also allows for interesting nonlinear interactions that we might suspect should be present. Consider the expected transaction cost and the logged value and volatility variables. We have $E(TC\%) = \exp^{(\text{other variables})} \text{value}^{\beta_1} \text{volatility}^{\beta_2}$. If β_1 is larger than 1 then the cost increases more than proportionally to the value. If β_1 is smaller than 1 then the expected cost increases less than proportionally to the value. If β_1 and β_2 are both positive, then increases in the volatility result in larger increases in the expected cost for larger value trades. Alternatively, as the value of the trade goes to zero, so does the expected cost. It is entirely possible that the marginal impact of volatility might be different for different order sizes. The exponential model allows for this possibility in a very parsimonious fashion. Hence, what appears as a very simple nonlinear transformation allows for fairly rich nonlinear interactions. Obviously, using the exponential function for the variance has the same interpretation.

We estimate the model by maximum likelihood under the normality assumption for ε . It is well known that the normality assumption is a quasi maximum likelihood estimator. As long as the conditional mean and variance are correctly specified, we still obtain consistent estimates of the parameters even if the normality assumption is not correct. The standard errors, however, will not be correct in the event that the errors are not normal. Robust standard errors that are consistent in the event of non-normal errors can be constructed following White (1982) and are constructed for our parameter estimates.

The estimation is performed separately for NYSE and NASDAQ stocks. The two markets operate in a very different fashion and it is unlikely a single model would be appropriate for both trading venues. We have 166,508 NYSE orders and 67,405 Nasdaq orders. The parameter estimates for the variance equation for the NYSE stocks is given in table 3.

The coefficient on the spread is positive indicating that wider spreads are associated with more risk for any given order type and order size. The coefficient on the log volatility is 1.2. A simple model where a given order type is always executed over the same time interval with roughly constant quantities traded implies that the variance of the transaction cost should be proportional to the variance of the traded asset. To see this, consider the variance of the transaction cost when the local effects are fixed so that the $\text{var}(\tilde{p}_t - p_t) = 0$. If equal quantities are traded in each time interval so that

$\frac{\Delta x_t}{(x_T - x_0)} = \frac{1}{T}$, and the variance of the asset is constant and given by σ^2 then the variance of the transaction cost:

$$(7) \quad \text{Var}(TC\%) = \text{Var}\left(\frac{\sum_{t=1}^T \Delta x_t (\tilde{p}_t - p_0)}{(x_T - x_0)P_0}\right) = \left(\frac{1}{T}\right)^2 \sum_{t=1}^T \text{Var}(\tilde{p}_t - p_0) = \left(\frac{1}{T}\right)^2 \sum_{t=1}^T t\sigma^2 = \sigma^2 \left(\frac{T^2 + T}{2T^2}\right)$$

For large T this is approximately $\frac{\sigma^2}{2}$ but the variance of the transaction costs should be proportional to the variance of the asset even for small T. Recall that $\text{Var}(TC\%) = \exp^{(\text{other variables})} \text{value}^{\beta_1} \text{volatility}^{\beta_2}$ so that it is therefore interesting to compare the estimated coefficient to the value 1. Squaring the volatility to convert the standard deviations to the variance $(\text{volatility}^2)^{\frac{\beta_2}{2}}$ yields a coefficient on the variance that is half the coefficient on the standard deviation which is .6 for the NYSE data. The variance of the transaction cost therefore increases less than proportionally to the variance of the asset. Thus, the Morgan Stanley BXS algorithm reduces the risk of the order relative to the simple constant volume, constant time interval strategy. This could happen for a

number of reasons including front loading the trades, or more rapid execution in higher volatility markets.

The coefficient on the log of the average 21 day median volume is $-.51$. Every 1% increase in the volume translates into a half a percent decrease in the trading cost. The order size has a coefficient of $.53$ indicating larger orders have a higher risk. A 1% increase in the order size translates into about a half of a percent increase in the variance. It is interesting to notice that the coefficient on the order size is roughly the negative of the coefficient on the volume. This indicates that logarithm of the order size as a fraction of the daily volume that predicts the variance. Not surprisingly, the variance of the transaction cost is decreasing as the urgency increases. This is consistent with the high urgency orders executing more quickly than the low urgency orders.

Next we turn to the mean cost parameter estimates. The spread is positively related to the transaction cost. A 1% increase in the spread translates into about a 1% increase in the transaction cost. Recall that the transaction costs are already expressed as a percent so this is a percent increase in the percent transaction cost. Wider spreads are consistent with markets that are less liquid. The volatility has a coefficient of $.50$. Every 1% increase in the 21 day volatility translates into a half of a percent increase in the expected trading cost. High volatility is often thought to be associated more uncertainty and less liquid markets as we find here. The coefficient on the average 21 day median volume is $-.47$. Every 1% increase in the daily volume translates into about a half of a percent decrease in the expected trading costs. The greater the volume the more liquid is the market.

The value has a coefficient of $.43$ indicating that a 1% increase in the value of the order translates into a little less than a half of a percent increase in the trading cost. It is again interesting to note that the coefficient on the value is roughly the same magnitude, but opposite sign as the coefficient on the volume. It appears that the size of the trade relative to the daily volume that predicts the cost. Finally, the cost is strictly increasing as the urgency increases.

The variance and mean model estimates for NASDAQ are presented in tables 5 and 6 respectively. While the magnitude of some of the estimates differs across the two exchanges the results are qualitatively very similar. We test the null hypothesis that the mean and variance models for the NYSE and NASDAQ are not different. This null hypothesis can be tested by an likelihood ratio test based on the difference between the sum of the likelihoods for the two unrestricted NYSE and NASDAQ models and the restricted model using the pooled data. Twice the difference in these two likelihoods will have a chi-squared distribution with degrees of freedom given by the number of restricted parameters, or 16. Twice the difference in the two likelihoods is 1954.76. The critical value is 26.29 so we overwhelmingly reject the null with a p-value near 0. Hence, while the models are qualitatively similar, there are statistically meaningful quantitative differences.

The parameter estimates provide intuitive interpretations regarding the transaction costs. It is nevertheless interesting to evaluate the statistical fit of the assumed exponential form. Toward this end we consider a variety of lagrange multiplier tests. The test can be performed for both omitted terms in the mean and the variance equations. Our null is that the exponential specification is sufficient while under the alternative we consider omitted linear and squared terms $E(TC\%_i) = \exp(X_i\beta) + Z_i\theta$ where Z will be taken to be X and X^2 or a combination of linear and squared terms. The test for the mean is performed by regressing the standardized error term on potential omitted terms. The standardized error term is given by $\hat{\varepsilon}_i = \frac{TC\%_i - E(TC\%_i)}{sd(TC\%_i)}$. We regress $\hat{\varepsilon}_i = \exp(X_i\hat{\beta})X_i\theta_0 + Z_i\theta_1$ where

X_i^2 is taken to mean the element by element square of each variable (ie no cross products are included). The θ_1 and θ_2 are conforming parameter vectors. Similarly, the test for the variance is performed by regressing $\hat{\varepsilon}_i^2 = \exp(X_i\hat{\beta})X_i\phi_0 + Z_i\phi_1$ where the ϕ_0 and ϕ_1 are again conforming parameter vectors. The results of these test and the special cases of omitted linear terms only and omitted squared terms only are presented in table 7. (TABLE 7 IS NOT READY YET).

Generally speaking, larger orders are cost more to execute than smaller orders. We next look more closely at how the expected cost and risk vary as the order size increases. Figures 3 and 4 plot the expected cost and the standard deviation as a function of the order size relative to the 21 day average median volume. The plots consider orders ranging from near 0 percent up to 2% of the daily volume. The plots are done for an average stock on an average day. Figures 5 and 6 present the same plots, but for the NASDAQ stocks. In the expected cost plots, the higher curves correspond to the more urgent orders. The opposite is true for the standard deviation plots.

We can also look at the conditional risk/cost trade off by plotting the mean and volatility conditional upon the state of the market for each order type. We again consider the risk/cost tradeoff for an average stock under average conditions. These contours are plotted in figures 7 and 8. The ellipses represent 95% confidence intervals for the true mean and true variance for each order submission strategy. As we move from left to right we move from high urgency to medium, to low and finally VWAP or the risk neutral strategy. Perhaps the most interesting conclusion is that this analysis suggests that there is not much benefit to moving from low urgency to VWAP for either NYSE or NASDAQ stocks. The change in the expected cost is nearly zero while the increase in risk is substantial. If the agent cares at all about risk, the VWAP strategy does not appear viable.

Given the model, we can evaluate the cost/risk tradeoff under any stock. To get an idea of how this tradeoff varies as we examine how the frontier changes as we vary the order size for a typical stock on a typical day. Again, it is natural to express the order size relative to the average median 21 day volume. These plots are presented in figures 9 and 10 for typical NYSE and NASDAQ stocks respectively. The larger orders shift the cost/risk tradeoff to less desirable north east region. We again see that the order size effects on the cost/risk tradeoff are substantial.

5. Liquidation Value at Risk (LVAR)

Liquidation risk is the uncertainty about how much it costs to liquidate a position in a timely manner if the need should arise. Liquidation risk is important from both an asset management/risk perspective, as well as a more recent literature on asset pricing and liquidity (see for example Easley and O'Hara (2003), Pastor and Stambaugh (2003), Pedersen and Acharya (2005)). The conditional distribution of transaction costs is fundamentally related to liquidation risk. We show how the losses associated with liquidating an asset can be bounded with some probability. We call this measure liquidation value at risk or LVAR. Like the traditional value at risk (VaR), LVAR tells us the minimum number of dollars that will be lost with some probability α , when liquidating an asset.

For a given liquidation order the conditional mean and variance can be constructed. Under a normality assumption one can construct an $\alpha\%$ LVAR given by:

$$(8) \quad LVAR(\alpha) = \exp(X\hat{\beta}) + \exp\left(\frac{1}{2}X\hat{\gamma}\right)z_{1-\alpha}$$

where $z_{1-\alpha}$ is the $1-\alpha$ % quantile.

More generally, we might not wish to impose the normality assumption and instead use a more non-parametric approach. In the first stage, consistent estimates of the parameters can be estimated by QMLE. In the second stage, the standardized residuals can be used to construct a non-parametric estimate of the density function of the errors ε . The standardized residuals are given by:

$$(9) \quad \hat{\varepsilon}_i = \frac{TC\%_i - \exp(X_i\hat{\beta})}{\exp\left(\frac{1}{2}X_i\hat{\gamma}\right)}$$

A non-parametric estimate of the density or perhaps just the quantiles themselves can then be used to construct a semi-parametric LVAR. Specifically, let $\hat{\varepsilon}_{1-\alpha}$ denote a non-parametric estimate of the $\alpha\%$ quantile of the density function of the error term ε . Then

the semi-parametric $\alpha\%$ LVAR is obtained by replacing $z_{1-\alpha}$ with the non-parametric quantile $\hat{\varepsilon}_{1-\alpha}$:

$$(10) \quad LVAR(\alpha) = \exp(X\hat{\beta}) + \exp\left(\frac{1}{2}X\hat{\gamma}\right)\hat{\varepsilon}_{\alpha}$$

Figures 11 and 12 present the standardized residuals for the NYSE and NASDAQ models. The residuals are clearly non-normal. We use the empirical quantiles of the data to construct the LVAR. Figures 13 and 14 present the 1% LVAR associated with the high, medium and low urgency orders as well as the VWAP. The LVAR estimates are constructed for typical stocks on a typical day. The vertical axis is the transaction cost in basis points. As we move from left to right we move from LVAR to low urgency to the high urgency orders. The LVAR is given by the upper bar for each order type. The expected cost for each order type is given by the smaller bar in near the origin. The differences in the mean are small relative to the changes in the risk across the different order types. Since the risk dominates, the minimum LVAR order type here is given by the most aggressive strategy, the high urgency order. The 1% LVAR for this order type is just under around half a percent for NYSE and 1% for NASDAQ. For each order type the lower dashed line completes a 98% prediction interval.

6. Conclusion

This paper demonstrates that expected cost and risk components of transaction costs can be estimated from detailed transaction data. We show that we can construct a cost/risk tradeoff in the spirit of classical portfolio analysis. We find that the expected cost and risk components can be successfully modeled using an exponential specification for the mean and variance. Characteristics of the order and state of the market play a major role in determining the cost/risk tradeoff faced by the trader.

We provide an example of how this approach can be used to assess liquidation risk using the notion of liquidation value at risk (LVAR). This is, of course, only one approach that could be taken in assessing liquidation risk. More generally, we have the entire conditional distribution of transaction costs so there are potentially many approaches that one could take in assessing liquidation risk.

Finally, our data here consists of the transaction costs. Another direction to go would be to directly consider the raw transaction data set. In this way, we could better assess the dynamics of the price impact functions. For example, how large are the local vs. price impact effects?

Exchange	Side	Benchmark	Urgency	Weight	Count	Price	Spread	Volatility	Volume	Capitalization (000)	Order Value	Order Shares	Cost (BP)	StDev (BP)
				100%	233,913	\$ 45.07	0.09%	26%	1.59%	\$ 59,609,060	\$ 310,472	9,154	10.09	47.24
	B			62%	147,649	\$ 45.06	0.09%	26%	1.57%	\$ 58,137,900	\$ 302,812	8,946	10.77	47.17
	S			38%	86,264	\$ 45.09	0.08%	26%	1.62%	\$ 61,965,453	\$ 323,583	9,512	8.99	47.31
NYSE				75%	166,508	\$ 48.01	0.09%	23%	1.68%	\$ 66,717,110	\$ 326,031	8,701	8.82	43.28
NASDAQ				25%	67,405	\$ 36.38	0.08%	36%	1.33%	\$ 38,565,201	\$ 272,037	10,273	13.84	57.19
		A	H	10%	15,616	\$ 47.81	0.08%	26%	1.18%	\$ 60,838,460	\$ 475,462	12,845	11.69	23.19
		A	M	24%	54,095	\$ 44.73	0.09%	27%	1.47%	\$ 48,482,781	\$ 320,909	9,688	11.09	32.20
		A	L	20%	51,588	\$ 46.44	0.08%	26%	1.13%	\$ 68,894,693	\$ 285,018	8,106	8.99	40.89
		V		46%	112,614	\$ 44.04	0.09%	26%	1.95%	\$ 61,042,206	\$ 294,240	8,867	9.69	59.01

Table 1. Summary statistics for Morgan Stanley trades. B and S are buy and sell orders respectively. A denotes arrival price strategy and V denotes VWAP strategy. H, M, and L denote high medium and low urgency trades.

Exchange	Benchmark	Urgency	Volume Range	Weight	Count	Price	Spread	Volatility	Volume	Capitalization (000)	Order Value	Cost (BP)	StDev (BP)
NYSE	A	H	≤ 0.25%	0.69%	2,630	\$ 47.77	0.08%	22%	0.19%	\$ 94,061,545	\$ 190,107	4.22	10.97
NYSE	A	M		3.10%	13,664	\$ 48.06	0.08%	22%	0.16%	\$ 95,714,257	\$ 164,767	3.69	11.64
NYSE	A	L		4.17%	19,379	\$ 50.15	0.08%	22%	0.14%	\$ 98,487,506	\$ 156,371	2.71	12.74
NYSE	V			6.54%	38,116	\$ 47.46	0.08%	22%	0.13%	\$ 90,733,082	\$ 124,606	1.97	34.56
NYSE	A	H	≤ 0.5%	1.69%	3,559	\$ 51.00	0.08%	22%	0.37%	\$ 81,476,811	\$ 345,605	6.16	11.95
NYSE	A	M		3.30%	9,557	\$ 49.09	0.08%	23%	0.36%	\$ 68,991,562	\$ 250,562	5.68	15.53
NYSE	A	L		2.99%	8,027	\$ 50.18	0.08%	21%	0.36%	\$ 92,105,328	\$ 270,537	4.15	19.65
NYSE	V			5.00%	14,890	\$ 48.23	0.08%	23%	0.37%	\$ 65,205,449	\$ 244,088	3.06	42.27
NYSE	A	H	≤ 1.0%	2.39%	2,979	\$ 52.69	0.08%	22%	0.73%	\$ 72,154,182	\$ 582,738	8.93	15.98
NYSE	A	M		3.64%	6,907	\$ 48.85	0.08%	24%	0.72%	\$ 54,822,342	\$ 383,195	7.54	20.67
NYSE	A	L		3.17%	6,035	\$ 50.22	0.08%	22%	0.72%	\$ 82,177,737	\$ 381,487	6.76	28.64
NYSE	V			6.40%	11,822	\$ 46.88	0.09%	23%	0.73%	\$ 71,181,217	\$ 393,015	5.17	47.42
NYSE	A	H	> 1.0%	2.67%	2,549	\$ 51.53	0.09%	23%	2.31%	\$ 48,375,509	\$ 760,186	14.36	25.12
NYSE	A	M		7.34%	7,052	\$ 47.41	0.10%	24%	3.02%	\$ 36,941,147	\$ 755,825	15.55	38.97
NYSE	A	L		4.77%	5,632	\$ 47.78	0.10%	23%	2.55%	\$ 52,146,046	\$ 615,387	12.64	52.73
NYSE	V			16.88%	13,710	\$ 45.67	0.09%	23%	3.95%	\$ 54,313,103	\$ 894,247	14.66	65.61
NASDAQ	A	H	≤ 0.25%	0.28%	715	\$ 36.70	0.06%	33%	0.18%	\$ 76,072,230	\$ 288,664	6.58	12.91
NASDAQ	A	M		1.59%	5,771	\$ 37.81	0.06%	34%	0.15%	\$ 52,884,846	\$ 200,367	6.05	17.13
NASDAQ	A	L		1.36%	5,065	\$ 37.49	0.06%	33%	0.14%	\$ 77,185,322	\$ 195,338	5.29	17.93
NASDAQ	V			3.73%	16,227	\$ 39.02	0.06%	34%	0.11%	\$ 71,786,694	\$ 167,134	3.96	42.05
NASDAQ	A	H	≤ 0.5%	0.82%	1,131	\$ 39.07	0.06%	32%	0.37%	\$ 63,042,552	\$ 527,242	10.76	18.02
NASDAQ	A	M		1.35%	4,455	\$ 36.72	0.08%	36%	0.36%	\$ 30,091,409	\$ 220,816	9.15	21.53
NASDAQ	A	L		0.89%	2,578	\$ 38.02	0.07%	36%	0.37%	\$ 37,310,972	\$ 251,330	8.33	31.04
NASDAQ	V			1.64%	5,597	\$ 35.82	0.07%	36%	0.36%	\$ 45,881,902	\$ 213,288	6.07	59.83
NASDAQ	A	H	≤ 1.0%	0.85%	992	\$ 38.30	0.07%	35%	0.70%	\$ 42,108,883	\$ 623,028	15.81	21.17
NASDAQ	A	M		1.31%	3,453	\$ 36.58	0.09%	38%	0.71%	\$ 15,709,279	\$ 275,812	14.78	30.01
NASDAQ	A	L		0.97%	2,003	\$ 39.11	0.07%	37%	0.72%	\$ 32,386,856	\$ 352,916	12.36	45.50
NASDAQ	V			1.80%	5,310	\$ 34.08	0.08%	36%	0.73%	\$ 36,671,191	\$ 246,208	10.62	66.66
NASDAQ	A	H	> 1.0%	0.83%	1,061	\$ 37.41	0.10%	36%	2.91%	\$ 10,249,188	\$ 565,871	26.99	47.16
NASDAQ	A	M		2.26%	3,236	\$ 32.85	0.11%	39%	3.10%	\$ 8,045,030	\$ 508,137	22.96	58.51
NASDAQ	A	L		1.91%	2,869	\$ 36.91	0.10%	38%	2.85%	\$ 15,236,379	\$ 484,223	26.02	79.01
NASDAQ	V			3.62%	6,942	\$ 33.30	0.11%	37%	3.46%	\$ 23,073,584	\$ 379,150	24.65	94.67

Table 2. Summary statistics for Morgan Stanley trades. Volume is the order size as a percent of the average daily volume. A denotes arrival price strategy and V denotes VWAP strategy. H, M, and L denote high medium and low urgency trades.

VARIABLE	COEFFICIENT	ROBUST T-STAT
Const	11.80559	90.86393
Spread	1.815802	14.39896
Log volatility	1.207152	64.98954
Log volume	-0.51614	-55.6044
Log value	0.536306	46.08766
Low urg	-1.45436	-83.2275
Med urg	-1.92541	-61.058
High urg	-2.33731	-88.2665

Table 3. Variance parameter estimates for NYSE stocks.

VARIABLE	COEFFICIENT	ROBUST T-STAT
Const	5.173827	30.0342
Spread	0.969804	8.395586
Log volatility	0.503987	21.14475
Log volume	-0.47084	-43.4163
Log value	0.43783	40.27979
Low urg	0.094929	2.41284
Med urg	0.305623	8.796438
High urg	0.41034	11.28093

Table 4. Mean parameter estimates for NYSE stocks.

VARIABLE	COEFFICIENT	ROBUST T-STAT
Const	11.40519	71.44173
Spread	2.016026	17.00666
Log volatility	1.078963	40.27656
Log volume	-0.44182	-46.7497
Log value	0.453704	42.64865
Low urg	-1.04398	-42.1013
Med urg	-1.70511	-65.0131
High urg	-2.10623	-48.7398

Table 5. Variance parameter estimates for NASDAQ stocks.

VARIABLE	COEFFICIENT	ROBUST T-STAT
Const	5.354067	26.04098
Spread	1.014023	8.734035
Log volatility	0.513628	16.96502
Log volume	-0.41447	-29.9304
Log value	0.376208	24.99588
Low urg	0.025356	0.243943
Med urg	0.230764	5.716912
High urg	0.282479	6.156517

Table 6. Mean parameter estimates for NASDAQ stocks.




Figure 1: NYSE average cost/risk tradeoff given the order size. The order size is expressed as a fraction of the median 21 day daily volume.




Figure 2: NASDAQ average cost/risk tradeoff given the order size. The order size is expressed as a fraction of the median 21 day daily volume.




Figure 3: Expected Cost as a function of the order size expressed as a fraction of average daily volume for NYSE stocks.




Figure 4: Standard deviation of transaction cost as a function of the order size expressed as a fraction of average daily volume for NYSE stocks.




Figure 5. Expected Cost as a function of the order size expressed as a fraction of average daily volume for NASDAQ stocks.




Figure 6. Standard deviation of transaction cost as a function of the order size expressed as a fraction of average daily volume for NASDAQ stocks.




Figure 7. Expected cost and risk frontier for a typical NYSE stock on a typical day.




Figure 8. Expected cost and risk frontier for a typical NASDAQ stock on a typical day.




Figure 9. Expected cost/risk frontier for a typical NYSE stock on a typical day. Each contour represents the frontier for a different quantile of order size expressed as a fraction of average daily volume.




Figure 10. Expected cost/risk frontier for a typical NASDAQ stock on a typical day. Each contour represents the frontier for a different quantile of order size expressed as a fraction of average daily volume.




Figure 11. Standardized residuals for NYSE stocks.




Figure 12. Standardized residuals for NASDAQ stocks.




Figure 13. This plot shows the 98% predictive interval for the transaction cost for a typical NYSE stock on a typical day. 0 corresponds to VWAP, 1 to low urgency, 2 to medium urgency and 3 to high urgency. For each trade type, the upper bar denotes the 1% LVAR.




Figure 14. This plot shows the 98% predictive interval for the transaction cost for a typical NASDAQ stock on a typical day. 0 corresponds to VWAP, 1 to low urgency, 2 to medium urgency and 3 to high urgency. For each trade type, the upper bar denotes the 1% LVAR.

References

- Acharya, Viral and Lasse Pederson(2005), “Asset Pricing with Liquidity Risk” *Journal of Financial Economics*,vol 11, pp.375-410
- Almgren, Robert and Neil Chriss, (1999) “Value under Liquidation”, *Risk*, 12
- Almgren, Robert and Neil Chriss,(2000) “Optimal Execution of Portfolio Transactions,” *Journal of Risk*, 3, pp 5-39
- Almgren, Robert,(2003) “Optimal Execution with Nonlinear Impact Functions and Trading-enhanced Risk”, *Applied Mathematical Finance*, 10,pp1-18
- Bertsimas, Dimitris, and Andrew W. Lo, (1998) “Optimal Control of Execution Costs,” *Journal of Financial Markets*, 1, pp1-50
- Chan, Louis K. and Josef Lakonishok, (1995) “The Behavior of Stock Prices around Institutional Trades”, *Journal of Finance*, Vol 50, No. 4, pp1147-1174
- Easley, David, Soeren Hvidkjaer, and Maureen O’Hara, 2002, Is information risk a determinant of asset returns? *Journal of Finance* 57, 2185-2222.
- Engle, Robert, and Robert Ferstenberg, 2006, Execution Risk, NYU manuscript
- Grinold R. and R. Kahn(1999) *Active Portfolio Management* (2nd Edition) Chapter 16 pp473-475, McGraw-Hill
- Obizhaeva, Anna and Jiang Wang (2005) “Optimal Trading Strategy and Supply/Demand Dynamics” manuscript
- Pastor, Lubos, Robert Stambaugh, (2003), “Liquidity Risk and Expected Stock Returns”, *Journal of Political Economy* 111, 642–685.
- White, Halbert, 1982, Maximum Likelihood Estimation of Misspecified Models, *Econometrica*, 50, 1-25.