# ITG ACE® – Agency Cost Estimator: A Model Description

**Investment Technology Group, Inc.**

***Financial Engineering Group***

*October 31, 2007*

# Table of Contents

# 1. Introduction

Investment performance reflects both the investment strategy of the portfolio manager and the execution costs incurred while implementing the objectives of the investment strategy. Execution costs can be large, especially when compared to gross returns, and thus can affect performance significantly. Managing execution costs can make or break the success of a particular investment strategy. For institutional traders who trade large volumes, implicit costs, most importantly the price impact of trading, typically represent a significant portion of total execution costs. See, for example, Domowitz, Glen, and Madhavan (2002) for various definitions of costs along with discussions and analyses.

The importance of accurately measuring execution costs has grown in recent years due to fragmented liquidity in today's equity markets, algorithmic trading, direct market access, and structural and regulatory changes such as decimalization (implemented in 2001) and Reg NMS (implemented in 2007). Moreover, the recent demand of some legislators and fund share holder advocates for better disclosure of commissions and other execution costs increases their importance even further (see, for example, Teitelbaum (2003)). This makes the management of execution costs an important issue for institutional investors whose trades are large relative to average daily volume. An excellent source of articles about measuring and managing transaction costs can be found, for example, in Investment Guides (2001) and Investment Guides (2002).

ITG provides a variety of tools to investors to help them minimize their execution costs, and hence maximize their realized returns. ITG ACE® (Agency Cost Estimator) is a mathematical/econometric model that provides a pre-trade estimate of the price impact costs of a given order. ITG ACE measures execution costs using the implementation shortfall approach introduced by Perold (1988), which defines execution costs as the appropriately signed difference between the average execution price and the prevailing price at the start of the order execution. This measure includes both the bid-ask spread as well as the price impact cost of the order – the two most important cost components. Explicit cost components, such as commissions, can easily be added to the ITG ACE estimate to obtain total costs of trading.

ITG ACE can be used in conjunction with other ITG pre-trade analytic tools to:

- provide accurate cost estimates (e.g., expected execution costs and standard deviation of execution costs of an order),
- estimate statistical characteristics of the distribution of execution costs, including distribution percentiles and confidence intervals,
- form pre-trade cost benchmarks to evaluate the execution performance of traders and brokers for a variety of common pre-specified strategies (in particular, VWAP-strategy – constant fraction of average daily volume, uniform strategy, ITG ACE Optimal Strategy) or any arbitrary user-specified strategy,
- analyze how the costs of trading depend on the trading strategy,
- fine-tune a trading strategy in terms of trading horizon, aggressiveness, and other parameters,
- find an optimal trading strategy that balances execution costs against the uncertainty in the realized costs of trading (opportunity costs).

In addition, ITG ACE can be used as post-trade cost benchmark for trading performance. Section 10 discusses *ITG Post-Trade ACE*, a generalized model that incorporates the general market effects at the time when the trading actually took place.

Unlike many other conventional products, ITG ACE is a dynamic model that recognizes that a trader or an automated system will typically need to break up a large order into several smaller trades to minimize price impact costs. There are three critical features of ITG ACE that merit special attention.

- ITG ACE recognizes that traders incur price impact costs because a trade moves the price adversely in the market when it is executed. It is the cost of demanding liquidity. Price impact has both a permanent and a temporary component. The *permanent component* is information-based: it captures the persistent price change as a result of the information the occurrence of a trade conveys to the market. The *temporary price impact* is transitory in nature: it is the additional price concession necessary to get the liquidity provider to take the other side of the order. The permanent price impact implies that the first trade of a multi-trade order will affect the prices of all subsequent sub-blocks sent to the market. Modeling this dynamic link is a key element of computing the price impact for a sequence of trades spread over time.
- ITG ACE recognizes that there is no such thing as “the” cost estimate of a trade. In reality, trading costs are a function of the trader’s strategy or execution approach. The more aggressive the trading strategy, the higher the costs. Trading aggressiveness can be measured in terms of

 

how rapidly the trader wants to execute the trade given the trade's size relative to normal volume. Thus, the ITG ACE estimate is based on a particular trading strategy.

- ITG ACE can also be used to find an "optimal strategy" that balances price impact costs against opportunity costs. Such an ITG ACE Optimal Strategy represents a solution of a very general optimization problem (with time-varying parameters) for both the single name and the portfolio case. It is a significant generalization of the solution originally introduced by Almgren and Chriss (2000). Opportunity costs are largely due to price volatility, which creates uncertainty in the realized costs of trading as it does for the realized returns of investing. When executing an agency order, the balance between price impact and opportunity costs is chosen on the basis of the motivation for the order, which is ultimately given by the investment manager. Passive managers are mainly concerned about price impact while growth or momentum managers are more worried about opportunity costs. We refer to the investment manager's sensitivity to opportunity costs as weight on risk, or risk aversion, just as is done for an investment manager's sensitivity to investment risk. The ITG ACE model estimates the expected costs and the standard deviation of the costs of the agency trading strategy that optimally balances the trade-off between paying price impact costs and incurring opportunity costs for a given level of risk aversion and trading horizon.[1] In ITG ACE, the user can define the weight on risk. To allow for this, the ITG ACE model formulates the trading problem as a multi-period stochastic control problem. The solution to this stochastic control problem is the optimal strategy that minimizes the weighted sum of price impact and opportunity costs. ITG ACE provides the expected costs and standard deviation of the costs for the resulting optimal strategy. This strategy is recommended for traders who want to weight the opportunity costs associated with trading over a long interval of time consistent with their weight on risk.

The ITG ACE model is not a purely econometric model calibrated based on transaction cost data. Rather, it is a structural model that uses parameters estimated econometrically. In particular, ITG ACE relies on stock-specific econometric models of volatility, price impact, and price improvement, as well as a risk model. In addition, a purely econometric model based on empirical data would not allow us to provide cost estimates for large orders since there simply are not many observations for large orders as

[1] The trading horizon can either be chosen by the user or ITG ACE can determine an optimal time horizon for a given order.

 

discussed in diBartolomeo (2006). By employing a structural model, ITG ACE does mitigate this problem.

It is very important to understand that the ITG ACE model is not a "magic button". Instead, it is a sophisticated tool to reliably forecast transaction costs and estimate their statistical characteristics for any scenario selected by a user. ITG ACE estimates depend on the user's strategy and the underlying price impact model parameters. The user's strategy is reflected in trading style and aggressiveness. The price impact model parameters are calibrated using proprietary ITG Peer Group data in order to be in line with "typical" costs of large institutions. The trading style is characterized by

- the aggressiveness (participation rate),
- the level of opportunistic trading.

The realized costs for opportunistic traders do not match with the costs of traders that have to execute most of the time. As a result, ITG ACE provides two cost estimates: one is called *ITG ACE Discretionary* and the other one is called *ITG ACE Non-Discretionary*.[2] As the names indicate, for ITG ACE Discretionary, we include all executions for the calibration, i.e., even orders for which the traders can postpone or abandon trading to take advantage of market conditions. For ITG ACE Non-Discretionary, we exclude opportunistic executions and only include orders for which the traders do not have much discretion and have to execute the orders no matter if market conditions are favorable or not. For more details, see Section 6.

Currently, ITG ACE is implemented for equities. However, there is increased interest in transaction cost models for non-equity asset classes. For example, Burghardt, Hanweck, and Lei (2006) develop a model with a risk-averse market maker that allows them to describe the shape of the limit order book. They apply their model to limit order book data for electronic futures contracts (E-minis).

[2] To be more specific, there are two versions of the ITG ACE model: ITG ACE/1 and ITG ACE/2. The ITG ACE/1 model is available for 42 countries and the ITG ACE/2 model is available for 21 of these countries. ITG ACE Discretionary and ITG ACE Non-Discretionary are available for ITG ACE/2 only, that is, for 21 countries. Please see Table 1 for the list of countries that are covered by ITG ACE/2 and also see the discussion of ITG ACE/1 and ITG ACE/2 in Section 6 of this paper.

 

ITG ACE is available for 42 countries. We estimate the ITG ACE model for each exchange of each country separately. This approach is necessary since transaction costs vary significantly between different countries and exchanges (see, for example, Munck (2006)[3]).

The next section describes the framework of the ITG ACE model.

# 2. Framework

The ITG ACE framework builds on the model introduced in Almgren and Chriss (2000). Similar model assumptions can also be found in Huberman and Stanzl (2005). The ITG ACE model is based on the concept of *bins*, i.e., the *trading horizon* is divided into periods of equal duration. For example for the U.S. market, ITG ACE considers thirteen 30-minute bins per trading day. In general, the model may be used for any number of bins of different duration as long as the parameters are estimated consistent with the duration. The trading horizon may consist of several days and the starting bin and the ending bin can be any bin on the first and last day, respectively.

An order is defined by its trading horizon, side (buy or sell), size, and *trading strategy* – i.e., the sequence of share quantities in each bin for a given trading horizon. It is assumed that trading of all share quantities is completed within their respective bins.

ITG ACE distinguishes between the *market price*, defined as a stock mid-quote price, and the *average execution price*, at which a given bin's shares are executed. The average execution price differs from the market price since it includes temporary price impact costs and average price improvement. For small orders this difference is typically only half of the prevailing bid-ask spread, net of any price improvement. *Price improvement* is defined as receiving a price better than the prevailing prices (bid for a sell or ask for a buy) at the time the order was placed. For larger orders that exceed the bid/ask size, the execution price reflects both permanent and temporary price impacts. *Permanent price impact* captures the information content of the order, while the *temporary price impact* is the cost of

[3] Munck (2006) also reports variation in transaction costs over time, which motivates why we update and re-estimate the ITG ACE model on a monthly basis.

 

demanding liquidity. Trade execution affects not only the trade price, but the market price as well. Large size trades move the market price not only within the execution period, but have a persistent effect on the market price to the end of the trading day. Such an effect is usually called a permanent price impact. The market price is also affected by other factors that are captured in a stochastic disturbance term. Of course, both the temporary price impact and the permanent price impact increase with the number of shares traded within a bin.

Finally, *execution costs* are the appropriately signed difference between the market price of the stock at the beginning of the trading horizon and the average execution price for the order. Since there are both deterministic and random factors involved in the dynamic analysis, execution costs are stochastic in nature and should be analyzed by statistical methods. Further, given the multi-period nature of the optimization control problem, the analysis also requires the use of stochastic dynamic programming.

Figure 1 provides an illustration to the above-described concepts and terms. The temporary and permanent price impact applies to both single and multiple executions.

Providing reliable estimates of the model's parameters presents a special challenge, and indeed is the most difficult aspect of creating and maintaining the ITG ACE model. Stock market dynamics are complex and are subject to a variety of institutional features. For example, price impact is extremely difficult to measure given the low signal-to-noise ratio induced by intraday price volatility, and very comprehensive statistical techniques to extract the "useful" signal are needed. In short, the econometric implementation of ITG ACE is the most critical element of the model development.

In the next sections, we turn to a discussion of these issues, focusing on the sub-models that underlie the ITG ACE parameter estimates. In particular, we review several theoretical price impact models published in the academic literature and compare them with the two price impact models used in the ITG ACE model.

# 3. Input Parameters

All ITG ACE implementations use stock-specific parameters estimated from the most recent market data, including security master information (ticker, cusip/sedol, exchange), the previous trading day's *closing price*, and estimates for *volatility*, *average trading volume*, and *bid-ask spread* of each security.

The volatility is the forward-looking price volatility from ITG's Risk Model.[4] Average trading volume is estimated as the median daily dollar volume for the 21 most recent trading days. The bid-ask spread is computed as the 5-day time-weighted average daily bid-ask spread. The estimation methodologies for average trading volume and bid-ask spread are selected to balance the latest trends in stock behavior against fluctuations generated by market news, earnings announcements, and other temporary factors. It is worthwhile to note that any other estimation approach can be used as well, if so desired.

The ITG ACE framework is built in such a way that the market price behavior of a stock may depend on its *expected intraday stock returns*. By default, these returns are set to zero, but client-specific "alpha" models may be included in the ITG ACE analysis.

Besides estimating transaction costs for single name trades, the ITG ACE model may also be used efficiently for pre- and post-trade analysis of portfolios. In all ITG ACE implementations, correlations between stock returns are estimated using ITG Risk Models (for details see the white paper "ITG Risk Models – Version 3" (2007)). Depending on the stock universe in the trade list, the corresponding country, region, or global ITG Risk Model is used.

The current implementation of the ITG ACE model supports 42 countries and there are plans to add more countries in the future.

# 4. Intraday Parameter Distributions

[4] Specifically, the ITG Risk Model Version 0 is used.

 

U.S. intraday volume and return patterns were originally studied by Harris (1986). He found that there are systematic intraday return patterns which are common to all of the weekdays, i.e., returns are large at the beginning and at the end of the trading day. The so-called U-shaped pattern of *intraday volume* and *volatility* has been documented since Harris' article in a number of other studies (see, e.g., Jain and Joh (1988) and Foster and Viswanathan (1993)).

McInish and Wood (1992), Lee, Mucklow and Ready (1993) and Chan, Chung and Johnson (1995b) study variation in the *intraday bid-ask spread* of New York Stock Exchange (NYSE) listed stocks. The authors find that the spread is widest at the opening, narrows during the day, and then widens again near the close. Intraday changes in inventory and information asymmetry risks are mentioned as possible explanations for the observed pattern. Chan, Christie and Schultz (1995a) and Chung and Van Ness (2001) consider intraday patterns of spreads for stocks traded on Nasdaq. In contrast to NYSE stocks, spreads of Nasdaq stocks decline throughout the entire day and the highest decline is at the end of the day. Chung and Van Ness (2001) show that intraday spread patterns of Nasdaq stocks changed significantly after SEC enacted major changes in the order handling rules on Nasdaq in 1997.

ITG ACE takes into account that trading volume, price volatility, and bid-ask spreads:

- vary significantly within the same trading day,
- change over the course of time,
- are stock-specific,
- are relatively stable for very liquid securities, and
- are not stable for illiquid securities.

The intraday variations in volume, volatility, and spreads are measured statistically and incorporated within ITG ACE's cost estimation. Ideally, if one intends to estimate costs for a stock, the intraday volume, volatility, and spread distributions for the particular stock should be used. The research, however, demonstrates that such distributions are unstable for less liquid stocks due to both market and stock-specific fluctuations. Figures 2 and 3 show intraday volume and spread distributions for stock XXX during several time periods.[5,6] Clearly, with such variation, for example, in the intraday volume

[5] Stock XXX is a randomly chosen mid-cap stock. The stock is a relatively illiquid stock, its market capitalization is about $390 million, and the median daily share volume is about 50,000 shares as.

or spread pattern for stock XXX, one cannot be certain that using the latest available distribution calculated from, e.g., March data will be a good estimate for April. A possible alternative for less liquid stocks is to use aggregated distributions based on a significant number of stocks, for example, all stocks included in similar markets (NYSE/AMEX, Nasdaq) and liquidity groups. These distributions are much more stable as demonstrated by the bold lines in Figures 2 and 3, and they provide more robust forecasts. We assume that distributions of trading volume, volatility, and spreads are, respectively, averages of trading volume, volatility, and spread distributions across individual stocks on an equally-weighted basis. All stocks included in this distribution are of equal importance. This makes sense, since the main purpose of the aggregation is to get meaningful and stable estimates for illiquid stocks. The same approach is applicable to international markets. Volume, volatility and spread distributions are updated monthly, based on the most recent available trade and quote data. Both stock-specific and aggregated distributions are smoothed to control for market noise.

ITG Financial Engineering produces a monthly report called "Intraday Parameter Distributions for World Equity Markets", which is distributed to ITG clients, containing aggregated intraday parameter distributions for 42 of the most liquid international markets. It demonstrates that intraday patterns vary considerably from country to country. For example, in contrast to the U.S. market, trading volume on the London Stock Exchange (LSE) is very low at the beginning of a trading day and increases throughout the day.

# 5. Trading Strategies

## 5.1 General Overview

In general, trading strategies can be subdivided into two categories: *structured* and *opportunistic trading strategies*.

[6] Note, in the remainder of this document, if not specified otherwise, all ITG ACE numbers presented in tables and figures are based on ITG ACE Non-Discretionary. For a description of ITG ACE Discretionary and Non-Discretionary, please refer to Section 6.2 of this document.

 

Opportunistic trading strategies do not strictly follow a pre-specified trading schedule. Instead, these strategies are continuously searching for liquidity and opportunities for favorable execution based on real-time information. The success of such algorithms requires reliable quantitative forecasts of price movements and liquidity patterns, as well as intelligently combined use of trading venues and alternative order types (such as discretionary limit orders, Immediate-Or-Cancel (IOC) orders, or pegging orders). Opportunistic trading strategies work well for orders that do not have to be completed. However, they are not suitable for orders that need to be executed in full within a certain time horizon.

In contrast, structured, or more precisely scheduled, strategies are generally linked to a certain benchmark, for instance Volume Weighted Average Price (VWAP) or implementation shortfall, and are mostly based on historical data and their underlying analytics like the historical intra-day volume, volatility, and spread patterns. At the macro-level, these algorithmic trading strategies suggest how to optimally slice a large order in different time intervals within a specified horizon, but additional intelligent rules have to be used to execute each part of the original order, taking specifically into account:

- how closely one should follow the suggested trading schedule (order timing, deviation rule),
- order type selection (limit orders, market orders, discretionary orders, and IOC orders, etc.),
- trading venue selection (smart order routing to execute at the best available price and to discover undisclosed liquidity).

Most of the rules require the input of real-time information and depend on models/algorithms that can be used to search for the best price with the fewest time constraints. For more information about strategy classifications and selections given the specific objectives and scenarios, see for example, Domowitz and Yegerman (2005) or Yang and Jiu (2006).

ITG ACE uses trading strategies that belong to the class of structured strategies. In ITG ACE, a strategy is defined as a sequence of number of shares that might best be executed within an execution period according to a bin scheme. A bin is a 30-minute period during a trading day. For example, in the U.S., 9:30-10:00 a.m. is bin 1 of day 1, 10:00-10:30 a.m. is bin 2 of day 1, ... , 3:30-4:00 p.m. is bin 13 of day 1; for multi-day strategies, 9:30-10:00 a.m. is bin 1 of day 2, etc.

There are several standard strategies that can be expressed by the bin scheme of ITG ACE:

- The *Instant Strategy* trades all shares in the starting bin. This strategy can be invoked in ITG ACE by setting any of the other strategies supported by ACE to start and end in the same bin.
- The *Uniform Strategy* assumes the same number of shares to be executed for each bin within the trading horizon. For example, if the order size is 300,000 shares and the trade should be completed between 10:00 a.m. and 1:00 p.m., the uniform strategy suggests executing 50,000 shares within each bin (bins 2 to 7). Bertsimas and Lo (1998) propose uniform strategies to minimize expected costs of trading fixed number of shares.
- The *VWAP Strategy by Horizon*. For each order input, ITG ACE generates a prediction of the stock's volume pattern over the desired time horizon, whether partial day, full day, or multi-day. For each order, the VWAP Strategy by Horizon is a trading strategy that matches the volume pattern of the underlying stock over the desired time horizon, participating more heavily during the periods when volume is expected to be heaviest. This helps to minimize the impact of trading during thin volume periods and allows the order to benefit from the most liquid conditions. Figure 4 presents the VWAP Strategy by Horizon for a trade of 300,000 shares of stock YYY that executes between 10:00 a.m. and 1:00 p.m.[7] The VWAP Strategy by Horizon is compared to the Instant Strategy, Uniform Strategy, and VWAP Strategy by Participation Rate with 30% participation rate. 300,000 shares of stock YYY represent approximately 8.5% of average daily volume (ADV).[8]
- The *VWAP Strategy by Participation Rate* is defined similarly to the VWAP Strategy by Horizon. For each order, the trading strategy is formed using the volume pattern of the underlying stock by participating proportionately with the specified participation rate in the estimated day's volume. If the fraction of order size relative to the average daily trading volume is larger than the participation rate, a multi-day strategy with the same intraday stock-specific volume pattern for each day is employed. Figure 5 displays four VWAP Strategies by Participation Rate with different participation rates (5%, 10%, 20% and 30%) for buying 300,000 shares of stock YYY. The trading always begins at 10:00 a.m. (i.e., in bin 2). The plot shows that the higher the participation rate is, the shorter the time horizon and thus the more aggressive the strategy.

---

[7] Stock YYY is a randomly chosen large-cap stock. The stock is a relatively liquid stock; its market capitalization is over $70 billion and the median daily share volume is 3.5 million shares.

[8] ADV is the median daily dollar volume for the 21 most recent trading days.

 

- The *ITG ACE Optimal Strategy* represents a solution of a very general optimization problem (with time-varying parameters). It is a significant generalization of the solution originally introduced by Almgren and Chriss (2000). The ITG ACE model estimates the expected costs and the standard deviation of the costs of the agency trading strategy that optimally balances the trade-off between paying price impact costs and incurring opportunity costs (for a given level of risk aversion and trading horizon.)

In what follows, we discuss ITG ACE Optimal Strategies and their interpretations and consequences in more detail.

## 5.2 ITG ACE® Optimal Strategies

The crucial question facing traders is how to define and quantify trading objectives in order to implement them in an appropriate strategy. This question is non-trivial since common trading objectives often compete with each other and cannot be completely satisfied simultaneously. For example, a cost-minimizing strategy is not necessarily the ideal solution. A trader who minimizes costs by breaking up a trade over a very long time horizon faces risk from significant market movements. But conversely, trading aggressively to control risk implies "front-loading" the order and typically raises costs. Therefore, an optimal strategy should balance both costs and risk. From this perspective, the ITG ACE Optimal Strategy is a valuable trading tool because it provides a mathematically derived optimal solution given certain model assumptions. These assumptions are discussed in detail below.

Execution costs are subject to a large number of unknown factors. These include, for example, the uncertainty caused by the behavior of other market participants and market movements related to macroeconomic or stock-specific factors. It is impossible to model all these factors. Therefore, we consider execution costs as a random variable rather than as a deterministic value or number. In other words, the same strategy may provide different results if it is executed repeatedly under the same circumstances. Generally, a probability distribution is characterized by a number of parameters. In particular, the mean and standard deviation are widely used in statistics as such parameters. Note that these parameters, in general, do not define a distribution uniquely, but if one assumes certain distributions, it is sufficient to consider only these two parameters to identify the distribution.[9] The mean of the distribution of costs may be interpreted simply as the average value of costs if the execution could be repeated many times. The standard deviation of costs characterizes how much the value of costs may deviate from the expected costs. Therefore, selecting a strategy best suited for given trading objectives is equivalent to selecting the best suited distribution of costs.

Clearly, every trader prefers both lower expected costs and lower risk (standard deviation of costs). Hence, both of these parameters enter the optimization objective function. To find the optimal trading strategy, we need to balance the trade-off between expected costs and the variance of costs. This yields the ITG ACE optimization problem

$$(1-\lambda)\cdot E(C)+\lambda\cdot Var(C)\to \min, \tag{1}$$

where $C$ is the *total execution costs* of the trade, *E(C)* is the *expected value of C*, and *Var(C)* is the *variance of C*. $\lambda$ is the ***risk** aversion parameter* in the interval [0,1]. $\lambda$ can also be considered as "weight on risk." The optimal solution is the trading strategy, among all strategies for a given set of trade side, trade size, and trading horizon, that minimizes the objective function in (1).

The *ITG ACE Optimal Strategy* is the solution of the optimization problem in (1). It is very important to realize that the solution depends on the trade characteristics and the selected risk aversion parameter. Different trade characteristics and different values of risk aversion produce different ITG ACE Optimal Strategies. Therefore, it is crucial to understand how to select the inputs into the optimization problem according to each particular situation.

## 5.3 Selection of Risk Aversion Parameter

The side and size of a trade are usually given, but a user may select the trade horizon and the risk aversion parameter. In order to select them more effectively, it is useful to remind ourselves that more

[9] The normal distribution is one widely used example of such distributions. Section 8 will discuss another example of such distributions

 

aggressive trading strategies have higher expected costs, but a lower standard deviation of costs. Both a shorter trading horizon and a higher value of risk aversion correspond to a more aggressive trading strategy. Figure 6 shows several probability distributions of execution costs for different risk aversions with a fixed one-day horizon for an order to buy 300,000 shares of stock YYY. The plot reveals that a higher risk aversion provides lower expected costs but higher standard deviation and thus, greater uncertainty. Therefore, a user should make a selection based on appropriate values of both expected costs and standard deviation of costs.

The following example demonstrates how one may make such a selection: Suppose we need to buy again 300,000 shares of the stock YYY in one day. We could trade the order using a variety of strategies – some more passive and some more aggressive. Each of these strategies has a corresponding risk aversion parameter. Figure 7 shows the possible expected cost/risk outcomes for various risk aversions. For most traders, a risk aversion of zero is too passive: while the expected costs are low, the risk is very high. The high risk due to the long trading horizon implies the possibility of executing at inferior prices – potentially destroying any alpha that a particular investment was anticipated to capture. However, if volatility in transaction costs is of no concern, then this strategy may be best since it will, over many orders, average to the lowest costs. Conversely, a risk aversion of one produces a very low-risk trading strategy, but with exceptionally high costs – yet another way to destroy alpha. The solution to avoiding these two extreme outcomes is to choose a risk aversion that appropriately balances costs and risk somewhere between the extremes.

Figure 7 shows that there are many choices of strategies between the extremes. As you move from left to right, you can see that you can incrementally reduce the expected costs of a trading strategy (relative to the most expensive) by assuming more risk. Somewhere along this “efficient frontier” of expected trading costs is a strategy that, beyond which, you begin to accumulate more risk than the reduction in expected costs is worth.

As Figure 7 demonstrates, trading strategies based on high risk aversion have low opportunity costs (opportunity costs are measured as the standard deviation of the transaction cost distribution). This lower standard deviation is achieved by trading more shares earlier in the trading horizon – which is closer to the decision price.[10] This “front-loading” tends to move the stock price more rapidly in the

[10] The decision price is the prevailing price at the time the decision to place the order is made.

unfavorable direction than an order executed more patiently. In the ITG ACE framework, this movement in the stock price is market impact. Therefore, if you desire low opportunity costs (low uncertainty in the transaction costs or low standard deviation) then you must be prepared to pay more market impact costs. If you are willing to keep open the chance of having large realized opportunity costs, you can slow the order execution down and avoid high market impact costs. Figure 8 shows optimal strategies for different risk aversions and fixed one-day horizons for ticker YYY and an order size of 300,000 shares. The chart shows that the larger the risk aversion the more front-loaded the strategy is. Figure 9 shows optimal strategies for risk aversion 0.3 and fixed one-day horizon for ticker YYY and for different order sizes. One can see in both figures that the risk aversion 0.3 (ITG ACE Neutral) yields a strategy that is close to a VWAP trading strategy. Moreover, the strategy becomes more and more back-loaded with increasing order size. This makes sense since market impact costs become more and more important.

Such a selection becomes more complicated if the trading horizon needs to be selected in addition to the risk aversion parameter, but the approach remains the same. As an alternative, ITG ACE can determine an "optimal" trading horizon for an order, thereby leaving the selection of risk aversion as the only user-specified input parameter.

## 5.4 Optimal Trading Horizons for Optimal Strategies

The selection of the trading horizon for an order is another parameter the user needs to choose. ITG ACE can provide an optimal trading horizon. The solution of finding such an optimal trading horizon may vary in practical situations. After considering client feedback and analyzing several different approaches, the following method generally proved to be the best suited for ITG ACE implementations. ITG ACE continues to increment the number of days by one until the expected transaction costs in equation (1) of the optimization problem decreases by less than a threshold value. In other words, the method suggests that there is no need to extend the trading horizon for one more day if the benefit of extending the horizon is not significant. This significance is determined by an algorithm that accounts for order size, costs, and volatility. The order-dependent threshold adjusts so that very large orders have a low cost to share value ratio as a threshold, whereas smaller orders have a higher cost to share ratio as a threshold. Additionally, more volatile names have a higher threshold since adding an additional trading day will increase the variance term significantly. In general, thresholds are around 3-5 bps but can be lower for very large order sizes. Table 2 shows expected costs, standard deviation of costs and optimal trading horizons for different values of risk aversion. The underlying orders are to buy 300,000 and 1,500,000 shares of the stock YYY. Expected costs and standard deviations are reported both in cents and basis points. The 300,000 and the 1,500,000 shares represent about 8.5% and 42.5% of ADV, respectively. The cost estimates are based on ITG ACE Optimal Strategies.

## 5.5 Incorporating Expected Returns and Optimization Constraints

Client-specific "alpha" models may be included into the ITG ACE analysis through input of intra-day expected returns. However, using non-zero expected returns to generate ITG ACE Optimal Strategies has one potential complication. ITG ACE may suggest optimal strategies which include orders of opposite direction to that of the overall order. For example, consider a sell order and assume constant positive expected intraday returns. As the stock price is expected to be higher at the end of the day, a profitable strategy for a trader is to buy shares at the beginning of the day and then sell the entire position at the end of the day at a higher price. Such a strategy is an optimal solution of the ITG ACE optimization problem, but users view it as undesirable since the strategy would try to benefit from short-term price movement predictions, which is not what ITG ACE is built for. ITG ACE can be constrained to require that all bin executions of the ITG ACE Optimal Strategy are on the same side of the market. Moreover, additional bin volume constraints can be added to the optimization problem such as trading at least 1% and at most 20% of historic average bin share volume in each bin.

# 6 Price Impact Model

## 6.1 Model Types

The modeling of both temporary and permanent price impact is the most complex and crucial part of ITG ACE. Various ways of specifying a price-impact function can be found in the academic literature. The simplest method is to assume a linear relationship between the (absolute or relative) price change caused by a trade and the trade's size. Typically, trade size is the number of shares executed, either in absolute terms or relative to the average (or median) total number of shares traded throughout the trade's duration.

Examples of articles that assume a linear price-order flow relation are Kyle (1985), Bertsimas and Lo (1998), Breen, Hodrick and Korajczyk (2002), and Farmer, J. D. (2002). Kyle presents one of the seminal market microstructure models that derives equilibrium security prices when traders have asymmetric information. In Bertsimas and Lo, the authors introduce a price impact model and apply stochastic dynamic programming to derive trading strategies that minimize the expected costs of executing a portfolio of securities over a fixed time period. Breen, Hodrick and Korajczyk develop a measure of liquidity and quantify the change in a stock price by the observed net trading volume. Farmer studies the internal dynamics of markets – for example, volatility clustering – proposing a non-equilibrium price formation rule.

Although initial models of price impact were linear with respect to trading volume, empirical evidence shows existence of non-linearities. Hasbrouck (1991a), (1991b) investigates non-linearities in the impact of trades on midquotes. He discovers an increasing, concave relation between price impact and order flow for several stocks traded on the NYSE. De Jong, Nijman and Roell (1995) use data on French stocks traded on the Paris Bourse and SEAQ International and show that the assumption of a linear impact of orders on prices is incorrect. Kempf and Korn (1999) use intraday data on German index futures to come to the same conclusion. Zhang (1999) offers a heuristic derivation of a non-linear market impact rule. For more discussions of empirical evidence concerning non-linearity of market impact, see e.g., Hausman, Lo and McKinlay (1992) or Chan and Lakonishok (1993). Nonlinear price impact models can be found, for instance, in Seppi (1990), Barclay and Warner (1993), Keim and Madhavan (1996), and Chen, Stanzl and Watanabe (2002). While Seppi, and Keim and Madhavan focus on the different impacts of block trades and market trades on prices, Barclay and Warner justify the non-linearity in the price-order flow relation by the "stealth-trading" hypothesis. This hypothesis claims that privately informed traders concentrate their trades in the medium size range. Since medium-size trades are associated with informed trading, larger trades add relatively little additional information. This results in a concave price-order flow relation.

ITG ACE supports two different price impact models, serving both the U.S. and international markets – **ITG ACE/1** and **ITG ACE/2**. Both methodologies belong to the non-linear class of models discussed above.[11]

ITG ACE/1 uses an enhanced version of the original ITG ACE price impact model. The original model assumed that price impact is a linear function of trade size, with coefficients based on stock-specific volume and volatility estimates. While this original version was only applicable for relatively small orders not higher than 30% of the stock's ADV, the enhanced ITG ACE/1 methodology provides meaningful transaction cost estimates beyond a 30% of ADV order size.

The ITG ACE/2 price impact model is a sophisticated mathematical/econometric model that is in line with recent academic empirical findings. It uses an econometric technique to estimate price impact functions based on market tick data. This technique is at the core of ITG ACE/2 and depends on several stock-specific parameters that are estimated daily and monthly using market data for every stock in the ITG ACE universe. Methods developed by the ITG Financial Engineering group have provided accurate estimates for different segments of the universe (exchange-specific and by liquidity group). This task is most challenging for illiquid stocks and varying methodologies are applied for segments of stocks with different liquidity characteristics. Permanent price impact coefficients are estimated based on one year's of tick data similar to the method in Hasbrouck and Seppi (2001). In particular, we aggregate trading for each stock over 30-minute intervals and measure price changes using the quote mid-points at the beginning and end of each interval. The observed price changes (normalized by the historical volatility for the bin) are regressed against the corresponding trade imbalances and approximated by a concave, bin-specific function. Assuming market equilibrium in the ACE framework, the resulting functions can be used to forecast the accumulated price impact within a 30-minute interval caused by partial fills of the order. Figures 10 through 13 show the empirical and theoretical ITG ACE permanent price impact functions for bin 1 for four different stock segments: the most liquid U.S. Listed stocks, all Listed stocks, the most liquid OTC stocks, and all OTC stocks. The

[11] This is a deviation from the framework in Almgren and Chriss (2000) who assume linear price impact functions. Almgren (2003) does allow for non-linear temporary price impact functions, while permanent price impact functions remain linear. ITG ACE allows for non-linear temporary and permanent price impact functions.

graphs show that the empirical functions become noisier when one restricts the stock universe. Nevertheless, all smoothed theoretical functions exhibit the same behavior and they can be characterized by three parameters: the slope *s*, the value *x* that represents the order size at which concavity starts and the concavity parameter *alpha*. Empirical evidence suggests that this behavior holds for all liquidity groups and all time intervals of the day. Figure 14 illustrates the intraday permanent price impact behavior for all segments of the Listed stock universe. The chart suggests that permanent price impact is the highest at the beginning of the day and lowest around noon and at the end of the day. Similarly to the intraday behavior of volume, volatility and spread, the intraday permanent price impact pattern is exchange- and country-specific. Figure 15 shows that the intraday permanent price impact behavior of stocks traded on Euronext is very different from the one in Figure 14 for U.S. Listed securities. Temporary price impact functions are estimated analogously and show very qualitatively identical properties.

Extensive research and testing with U.S. and international execution data have demonstrated the accuracy of the approach for orders up to 100% ADV (see Section 9). The ITG ACE/2 price impact methodology is available for the U.S. market and the most liquid international markets (21 countries in total). Table 1 lists the countries covered by ITG ACE/1 and ITG ACE/2.

## 6.2 Discretionary and Non-Discretionary Liquidity Demand

In ITG ACE/2, the magnitude of price impact for each security and order size is defined by a quarterly calibration to ITG's Peer Group Database (see Section 9 for more details). As such, the price impact functions are sensitive to the orders contained in the database. Since the database is extremely large and comprehensive, it contains executions representing not only a wide spectrum of sizes, brokers, execution venues, and stock characteristics, but a broad range of trading behavior stemming from investment management styles, market conditions, trade motivations, and news events. This richness of the dataset allows ITG the unique opportunity to provide transaction cost estimates that reflect more than a "market average" trading behavior.

 

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 move against the trader. 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 ITG's 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 ITG 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. From a pre-trade perspective, the *ITG 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 *ITG ACE Discretionary* cost estimate provides a number that is suitable for comparing incurred transaction costs with what other 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 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 (ITG 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 *ITG ACE Non-Discretionary* and *ITG ACE Discretionary* are derived from different subsets of the same peer group database. For the *ITG ACE Discretionary* model, the entire database of orders minus those eliminated by outlier filtering are included in the calibration process. For the *ITG ACE Non-Discretionary* model, a sophisticated set of heuristics is used to eliminate database participants that exhibit opportunistic trading. These methods focus on identifying participants whose orders do not meet minimum transaction cost requirements with respect to increasing order size.

Figure 16 and Figure 17 plot the average realized costs curves that are associated with ITG ACE Discretionary and ITG ACE Non-Discretionary along with the average realized cost curve for opportunistic orders for Listed and OTC stocks, respectively. In both charts it is apparent that opportunistic orders are very different, they have very low costs, often close to zero and costs do not increase with order size. The cost curve associated with ITG ACE Non-Discretionary is above the cost curve associated to ITG ACE Discretionary, as expected. Excluding the opportunistic orders pushes the cost curve up. As discussed above, the difference in the curves is bigger the larger the order size is.[12]

Figure 18 and Figure 19 show the difference in cost estimates for stocks XXX and YYY using *ITG ACE Non-Discretionary* and *ITG ACE Discretionary* for various order sizes. In Figure 18, the cost estimates are based on a VWAP by Horizon Strategy with a one-day trading horizon. As expected, transaction costs for orders that need to be completed are higher than those that reflect a market average amount of opportunistic trading. In Figure 19, the cost estimates are based on a VWAP By Participation Strategy with 10% participation rate. As a result, orders can span multiple days. We observe the same pattern, transaction costs for orders that need to be completed are higher that that reflect a market average amount of opportunistic trading. Compared to Figure 18, the cost estimates are higher for very small orders, but lower for larger orders. The one-day horizon in Figure 18 forces the execution of an order into one day even if for larger sizes. This explains the higher costs in Figure 18 for larger orders compared to Figure 19. For very small orders, the logic works the other way around. Whereas the one-day horizon in Figure 18 allows for the order to be spread over the entire day, the 10% participation rate in Figure 19 forces the execution in the early half-hour intervals of the

[12] This makes intuitive sense: it is likely that the larger an order is, the more care that is applied and the more discretion that is given to the trader.

trading day. This leads to higher cost estimates for two reasons. First, the trading is concentrated in the early bins and at 10% participation rate may be much higher than the one-day horizon trading rate in Figure 18. Second, spread costs are highest early in the morning (see Figure 3), and thus the 10% participation strategy incurs those higher spread costs early in the morning. There is one more observation in Figure 19 that needs explanation. For stock XXX, the ITG ACE Non-Discretionary cost estimates are higher declining in order size for he very smallest order sizes. The explanation, again, is due to the fact, that for small orders, the 10% participation rate will imply full execution of the order in the early morning, thereby incurring the spread costs that are highest in the early morning. By increasing the order incrementally, cost estimates actually go down since the costs due to spread costs are declining as the order is spread more and more into the day outweighing any price impact costs that arise with larger order size. This effect subsides and the effect of larger price impact for larger orders takes over at a certain order size resulting in the usual increasing cost function. For ITG ACE Discretionary, we do not observe this pattern since opportunistic traders may use limit orders and time their trading such that the spread costs do not have an impact on their costs and the lower costs of the opportunistic traders outweighs the effect from the non-opportunistic traders.

# 7. Price Improvement Model

Generally, all buyer- (seller-) initiated orders are expected to be executed at the prevailing ask (bid) price. However, a trader may often achieve a better execution price and, therefore, realize a price improvement. Price improvement may appear simply because the market moved favorably during the time it took to route the order to the exchange, resulting in a lucky saving. But there are also other more sophisticated market microstructure theories why price improvement occurs. An excellent overview can be found in Rhodes-Kropf (2002). The discussion there is focused mostly on price improvement in dealership markets. Petersen and Fialkowski (1994) and Ready (1999) explain the existence of price improvement in auction type markets like the NYSE through hidden limit orders or stopped orders. For details about hidden limit orders and how to predict the volume executed against hidden limit orders for different market conditions, see e.g., Bongiovanni, Borkovec and Sinclair (2006).

The ITG ACE Price Improvement model allows users to quantify the price improvement of small size orders for different exchanges and values of order side, size, and liquidity. The model is based on ITG proprietary execution data for U.S. and the ITG Peer Group Database for international orders, respectively. These sources provide the necessary information to obtain market prices and to measure price improvement at any particular moment of trade execution. Not surprisingly, the results indicate that price improvement can be very different for quote- and order-driven stock markets.

We calculate relative price improvement for different exchanges, trade sides, trade sizes, and groups of liquidity as

$$R = \delta \cdot \frac{(p_Q - p)}{(p_{ask} - p_{bid})} \quad , \tag{4}$$

where $p$ is the trade price, $p_{bid}$ and $p_{ask}$ are the prevailing bid and ask quotes, respectively, $p_Q = p_{ask}$ and $\delta = 1$ for buys, $p_Q = p_{bid}$ and $\delta = -1$ for sells. Such a parameter has a very clear interpretation. For relatively small trades the value of $R$ usually lies between 0 and 0.5. If a buy (sell) trade was executed at the ask (bid) price, $R$ is equal to 0; i.e. there was no price improvement.

Figure 20 demonstrates that the average empirical relative price improvement for stocks traded on the NYSE depends on trade size and trade side. The graph is based ITG proprietary execution data for June 2006. Highest average relative price improvement occurs for the smallest trades and decreases as trade size increases.

Figure 21 compares average empirical relative price improvement for stocks traded on the NYSE that belong to different liquidity groups. The plot shows that there is almost a linear relation between average relative price improvement and liquidity. Relative price improvement is the lowest for the most liquid stocks and the highest for the most illiquid stocks. In addition, note that price improvement in absolute terms can be still higher for illiquid stocks due to the generally much larger spread. Sell trades, on average, obtain more price improvement than buy trades.

# 8. Distribution of Transaction Costs

As discussed in the previous sections, the execution of orders can be thought of as a trade-off between the risk of delayed execution and the cost of immediacy (see also, Hasbrouk and Schwartz (1988)). Much research has focused on the optimal execution of orders under various assumptions. Various forms of market impact models have been considered by practitioners, using theoretical or empirical methods to develop a set of market impact functions, both temporary and permanent (e.g., ITG ACE/1, ITG ACE/2, Kissell and Glantz (2003), or Almgren et al. (2003)). A common feature of these approaches is the assumption that the uncertainty in transaction costs can be represented entirely by the volatility of the security's return. The implication of this assumption is that there is no interplay between trading activity and a security's return volatility. This requires that the market in the security is near equilibrium during trading, that is, the security's return volatility remains constant while the return itself is affected by the market impact due to the trading. The assumption of independence of the moments of the return distribution and trading seems unrealistic. Almgren (2003) makes some important advances in the study of the interaction between trading activity and observed volatility. He derives optimal execution strategies for cases where volatility increases linearly with trading rate. ITG ACE takes a different approach, rather than modeling a security's return volatility conditional on trading, ITG ACE models the uncertainty in transaction costs directly as discussed below.

Typically, a portfolio manager will construct a portfolio on the basis of net returns (i.e., gross alpha less transaction costs). Such a model provides not only expected transaction costs, but also an uncertainty measure associated with it. Often, a moderately volatile stock will exhibit uncertainty of equal to or of an even much greater magnitude than the expected costs, so a good measure of the uncertainty in transaction costs resulting from the security's return volatility under liquidity pressure is crucial to an accurate transaction cost model. When analyzing *ex-post* trading performance, this same uncertainty about transaction cost estimates is used to determine the quality of execution. A trading desk manager may ask: "Did 67% of trading costs fall within one standard deviation of the expected trading costs?" Basing the answer to this question on a security's return volatility estimates, rather than the actual expected distribution of transaction costs will be misleading due to the described dependence between return volatility and trading.

A second concern is that most previous work on optimal trade execution has assumed constant, normal distributions of security returns during trading. We have shown that a large cross-section of actual executions[13] exhibit fat tails and skewness not accurately described by a normal distribution. Instead, we find that the distribution of transaction costs can be accurately modeled with an asymmetric generalized t-distribution. The generalized t-distribution was introduced by McDonald and Newey (1988) and the skewed extension of it was proposed by Theodossiou (1998). The family of asymmetric generalized t-distributions is very flexible and includes five parameters: two parameters $p$ and $q$ define the general shape of the distribution (Figure 22 illustrates some examples of generalized t-distributions with different choices of $p$ and $q$), one parameter $\alpha$ defines the asymmetry of the distribution and the final two parameters are location and scale parameters that determine the mean and variance of the distribution. The generalized asymmetric t-distribution contains many families of distributions, amongst them are the normal distributions ($p=2$ and $q \to \infty$) and the Student's t-distributions ($p=2$ and $q = 2\beta$, where $\beta$ denotes the degree of freedom of the Student t-distribution).

The ITG ACE transaction cost distributions are generalized asymmetric t-distributions with fixed, order-independent coefficients p, q and $\alpha$ while the location and scale parameters reflect the expected cost of the order and the security's return standard deviation over the trading horizon adjusted by the order size relative to the security's ADV. The adjustment is in line with Algren (2003) and empirical evidence that predicted standard deviations of transaction costs solely based on the security's return are lower than the empirical standard deviations. The adjustment of the standard deviation and the shape and asymmetry coefficients are derived from ITG Peer Group data similarly as described in Arellano-Valle et al. (2004). Figure 23 presents the fit of the empirical distribution of the z-scores of all actual costs with the ITG ACE z-score distribution (determined by the three parameters p, q and $\alpha$). For the purpose of illustration we have added some other calibrated theoretical distributions. Clearly, the generalized t-distribution outperforms all other distributions. Statistical techniques such as the Kolmogorov-Smirnoff test confirm this fact.

In summary, the ITG ACE cost distributions are characterized by three fixed parameters, the expected transaction costs and the standard deviation of the transaction costs. However, since the cost distributions are not normal distributions, one needs to use care when constructing confidence intervals

[13] From ITG's Peer Group Database.

based on mean and standard deviation. The usual interpretation that mean +/- one standard deviation contains two thirds of the observation no longer applies. Consequently, it is beneficial to also look at percentiles of the distribution. The percentiles of the cost distribution for a given scenario are part of the output of ITG ACE.

# 9. Calibration and Testing

As for any model, the key question for ITG ACE is how well the model actually performs. The accuracy of the model is controlled and validated through a process of calibration and statistical testing. The goal of calibration is to tune the price impact coefficients derived from market tick-data to achieve an alignment with realized transaction costs from a large database of known orders. Statistical testing is used to ensure that the model is returning unbiased results (i.e., costs that are not systematically over- or underestimated.)

## 9.1 Calibration

Each quarter, ITG ACE is calibrated to ITG's Peer Group Database. A moving two-year span of data is used, comprised of approximately four trillion U.S. dollars in trades from over 85 large investment management firms. For more information about the underlying data see Table 3 for the most important countries of the ACE universe.

Establishing a suitable data set for calibration and testing is a difficult endeavor for several reasons. Firstly, execution data often do not contain as much detailed information as is desirable. For example, execution and decision times might be missing or there is no clear declaration if the underlying order was a market or a limit order. Secondly, transaction costs depend on execution strategies and these strategies are, in most cases, not formalized by traders and certainly not recorded. Due to numerous factors (e.g., market conditions, work load, explicit instructions from portfolio managers) it is very likely that traders execute similar trades very differently over the course of a year.

Finally, an additional challenge exists in finding an approach to discount significant market and/or stock-specific movements, allowing for the measurement of the pure unperturbed magnitude of transaction costs. To this end, ITG carefully establishes a methodology that reflects the needs of the calibration and testing processes, while being sensitive to challenges presented by the data.

Since investment managers' orders are often broken into smaller orders or trades, an aggregation must be performed before arriving at a basic order unit suitable for analyzing trading activity, its effect on prices and, thus, comparison with ACE average cost estimates. To perform this aggregation, trade packages (ex-ante orders) are created that correspond to groups of trades where the same investment manager is in the market for a stock (buying or selling) over a sustained period of time.

The clusterization concept is in line with academic literature (see e.g., Chan and Lakonishok (1995)) as well as industry practice. The entire sequence of trades (ex-ante order) is treated as the basic unit of analysis in order to determine price impact and execution costs of institutional trading. In particular, a "buy ex-ante order" is defined to include the manager's successive purchases of the stock. The order ends when:

(a) the manager stays out of the market for at least one day,

(b) the manager does not execute more than 2% of ADV,

(c) there are no other trades that have been placed as an order within the execution horizon of the package.

"Sell ex ante orders" are defined analogously. For each ex-ante order, the trading aggressiveness (participation rate) and the average execution price is determined. Since execution time stamps are generally not reported, it is assumed that each ex-ante order has been executed according to a VWAP strategy with the empirically estimated participation rate. In most cases, this assumption is reasonable since large institutions are often measured against the VWAP benchmark.

The transaction costs per share are defined as the difference between the average execution price and the opening price of the order placement date (the benchmark price). The sign (positive or negative) of the difference is used so that a positive value represented a bad outcome. For each ex-ante order in the data set, the realized transaction costs $x_t$ are computed. Also calculated, using the parameters of the

ITG ACE model and the actual trading strategy for each order, are the estimated expected transaction costs $m_t$. This enables a one-to-one comparison between actual and estimated transaction costs.

For model calibration and testing, average actual costs and ITG ACE estimates are computed for the data set, segmented by size relative to ADV, by exchange, and by liquidity group. More specifically, for a given exchange and liquidity group, orders are subdivided into the following different size categories: 0-1%, 1-2%, 2-3%, … , 98-99%, 99-100%.

A two-step regression approach is applied to ensure that average actual costs and ITG ACE cost estimates coincide. Loosely speaking, the calibration procedure adjusts the price impact coefficients in such a way that the average ITG ACE cost estimates fit to the actual average costs. The adjustment is applied uniformly across all bins in order to avoid destroying the intra-day relationship of the price impact coefficients. As a consequence, low actual average costs will imply low price impact coefficients and therefore low ITG ACE cost estimates. Figures 24 to 29 serve as examples for the goodness-of-fit of empirical cost curves from the ITG Peer Group Database and the calibrated ITG ACE model.

The quality of the calibration depends crucially on the amount of data that is available. If a segment has not sufficient observations we merge data, more precisely, we combine ex-ante orders that belong to different liquidity groups. The calibration process is monitored for each market venue and two reports are generated: a *calibration report* and a *stability report*. The calibration report shows the goodness-of-fit, i.e., it compares average realized costs with average ITG ACE cost estimates for different order sizes (both equally- and local currency-weighted) for each primary exchange. The stability report summarizes the stability of the calibrated price impact functions for each liquidity group and primary exchange. If a huge deviation from one quarter to the other appears, a closer investigation is required. Large deviations can happen because of bad/corrupt data, not sufficient observations, or general changes for a market venue (e.g., trading hour changes).

## 9.2 Testing

On an annual basis, ITG ACE undergoes a battery of tests to help ensure that the model is free from statistical bias. Considering that the nature of transaction cost distributions are, as empirical evidence suggests, heavy-tailed or leptokurtic, simple parametric test procedures like the t-test are not suitable. Instead, two non-parametric statistical tests are employed to determine the accuracy of ITG ACE: the sign test and the Wilcoxon signed-rank test.

For the sign test, the null hypothesis is that the median difference between the ITG ACE cost estimate and the realized costs is zero. That is, after ranking the differences between the ITG ACE cost estimate and realized costs, the middle value should be zero. If much more observations are higher than (lower than) the actual value, the sample statistic can become too high and the sign test reject the null hypothesis.

For the Wilcoxon signed-rank test, the null hypothesis is that the differences between the ITG ACE cost estimates and realized costs are symmetric around zero. To calculate the Wilcoxon signed-rank test statistic, the differences from the sign test are used. However, the overestimates are separated from the underestimates. The sum of the ranking values for each group are calculated and compared. The statistic is the larger of the sum of the ranks of the over- or underestimates. If there were ten observations, the maximum value of this statistic is 55 (one through ten summed together).

The sign and Wilcoxon signed-rank test statistics are computed for the data set after it is segmented by size relative to ADV, exchange, and market capitalization. More specifically, orders are subdivided into different size categories using the following three different subdivisions. Note that, in the third subdivision, due to relatively few observations and wide dispersions of reported costs, test results for the group “greater than 50% of ADV” are less informative than others.

1) Uniform Subdivision 1:
   0-4%, 4-8%, 8-12%, 12-16%, 16-20%, and 20-24% of ADV
2) Uniform subdivision 2:
   0-15%, 15-30%, 30-45%, 45-60%, 60-75%, and 75-90% of ADV
3) ITG TCA subdivision (same as in ITG TCA, ITG’s product for post-trade analysis ):
   0-1%, 1-5%, 5-10%, 10-25%, 25-50%, greater than 50% of ADV

 

Next, orders in each size category (e.g., less than 10% of ADV) are subdivided into categories by exchange and market capitalization – for the U.S.: OTC, Listed, small-cap OTC, mid-cap OTC, large-cap OTC, small-cap Listed, mid-cap Listed and large-cap Listed. For each size category as well as the eight exchange and market capitalization groupings, the sign and Wilcoxon signed-rank test statistics are calculated, creating 96 tests of statistical significance of the accuracy of ITG ACE for the U.S. Based upon the distributions of the sign and Wilcoxon signed-rank tests, the results are checked to see if the calculated values are within the 95% confidence interval of the null hypotheses. When conducting a statistical test with 5% significance level, one expects the test to fail five times in 100 even if the null hypothesis is true. In the case for the U.S. with 96 tests, we would expect the tests to fail about five times. The results from the tests of ITG ACE Discretionary for the U.S., Canada, the U.K., Japan, and Euronext (consisting of the exchanges in France, Belgium, The Netherlands, and Portugal) for the time period January 2005 to December 2006 are provided in Tables 4 through 8.[14]

# 10. ITG ACE® For Post-Trade Analysis

Pre-trade transaction cost models typically measure the institutional average price impact costs. The embedded crucial assumption is market neutrality. Consequently, the estimated costs are entirely based on one's own trading strategy and the associated price impact. Pre-trade models per se do not account for market effects due to other market participants, short-term serial correlations in price movements,[15] news events/announcements, and the underlying investor sentiment in the market. The ITG ACE model is not an exception for the pre-trade usage. However, for post-trade analysis, incorporating general market effects such as actual returns and trade imbalances at the time when the executions happen is possible. Consequently, one can potentially significantly improve transaction cost estimates for post-trade analysis.

The ITG ACE For Post-Trade Analysis model builds exactly on this idea. It uses the ITG ACE framework augmented by the actual market conditions during the actual execution period of an order.

[14] Note, depending on the number of exchanges considered for each of the countries, the number of tests varies between 48 for Canada, 96 for the U.S., the U.K., and Japan, and 192 for Euronext. For details, please see Tables 4 through 8.

[15] Price movements exhibit serial correlation across various time horizons (Lo and MacKinlay (1988)) and reflect changes in market conditions and the presence of privately informed traders (Bertsimas and Lo (1998)).

 

The actual market conditions are summarized by factors such as market, sector, and industry returns as well as stock-specific characteristics such as bid-ask spread and trade imbalance information. Extensive tests have shown that ITG ACE For Post-Trade Analysis is a very reliable post-trade transaction cost benchmark. In addition, the model also allows for the decomposition of the costs of a transaction into two components: the costs due to one's own trading and the costs due to general market effects. For more details about ITG ACE For Post-Trade Analysis, see Borkovec and Heidle (2007).

# 11. References

Almgren, R. (2003) Optimal execution with nonlinear impact functions and trading-enhanced risk, *Applied Mathematical Finance*, **10**, 1-18.

Almgren, R. and Chriss, N. (2000) Optimal execution of portfolio transactions, *Journal of Risk*, **3**, Winter 2000/2001, 5-39.

Almgren, R., Thum, Ch., Hauptmann, E. And Li, H. (2003) Direct estimation of equity market impact, http://www.courant.nyu.edu/~almgren/papers/costestim.pdf.

Arellano-Valle, R. and Gómez, H. and Quintana, F. (2004). A New Class of Skew-Normal Distributions, *Communications in Statistics, Series A*, **33**, 1465-1480.

Barclay, M. J. and Warner, J. B. (1993) Stealth trading and volatility, *Journal of Financial Economics*, **34**, 281-305.

Berkowitz, S., Logue, D. and Noser E. (1988) The total cost of transactions on the NYSE, *Journal of Finance*, **41**, 97-112.

Bertsimas, D. and Lo, A. W. (1998) Optimal control of execution costs, *Journal of Financial Markets*, **1**, 1-50.

Bongiovanni, S., Borkovec, M. and Sinclair, R. (2006) Let’s play hide-and-seek: the location and size of undisclosed limit order volume, *Journal of Trading*, Summer 2006, 34-46.

Borkovec, M. and Heidle, H. (2007) The magic of hindsight: creating a post-trade transaction cost estimate based on realized market conditions, *ITG Inc. Research Report*.

Breen, W.J., Hodrick, L.S. and Korajczyk, R.A. (2002) Predicting equity liquidity, *Management Science*, INFORMS, **48**, 470-483.

Burghart, G., Hanweck, J. and Lei, L. (2006) Measuring Market Impact and Liquidity, *Journal of Trading*, Fall 2006, 70-84.

Chakravarty, S., Panchapagesan, V. and Wood R.A. (2002) Did decimalization hurt institutional investors? *Journal of Financial Markets*, **4**, 400-420.

Chan, L.K.C. and Lakonishok, J. (1993) Institutional trades and intraday stock price behavior, *Journal of Financial Economics*, **33**, 173-199.

Chan, L.K.C. and Lakonishok, J. (1995) The behavior of stock prices around institutional trades, *Journal of Finance*, **50**, 1147-1174.

Chan, K., Christie, W. and Schultz, P. (1995a) Market structure and the intraday pattern of bid-ask spreads for Nasdaq securities, *Journal of Business*, **68**, 35-60.

Chan, K., Chung, P. and Johnson, H. (1995b) The intraday behavior of bid-ask spreads for NYSE stocks and CBOE options, *Journal of Financial and Quantitative Analysis*, **30**, 329-346.

Chen, Z., Stanzl, W. and Watanabe, M. (2002) Price impact costs and the limit of arbitrage, *http://students.som.yale.edu/Phd/mw239/research/Price_Impact.pdf*

Chung, K. H. and Van Ness, R. A. (2001) Order handling rules, tick size, and the intraday pattern of bid-ask spreads for Nasdaq stocks, *Journal of Financial Markets*, **4**, 143-161.

De Jong, F., Nijman, T. E. and Roell, A. (1995) A comparison of the cost of trading French shares on the Paris Bourse and on SEAQ International, *European Economic Review*, **39**, 1277-1301.

diBartolomeo, D. (2006) Building a better bear trap – new concepts in cost and impact modeling, *Journal of Automated and Algorithmic Trading*, July 2006, 18-21

Domowitz, I., Glen, J. and Madhavan, A. (2002) Liquidity, volatility and equity trading costs across countries and over time, *International Finance*, **4**, 221-256.

Domowitz, I. and Yegerman, H. (2005) Measuring and interpreting the performance of broker algorithms, *ITG Inc. Research Report*.

Farmer, J. D. (2002) Market force, ecology, and evolution, *Industrial and Corporate Change*, **11**, 895-953.

Foster, F. D. and Viswanathan, S. (1993) The effect of public information and competition on trading volume and price volatility, *Review of Financial Studies*, **6**, 23-56.

Harris, L. (1986) How to profit from intradaily stock returns, *Journal of Portfolio Management*, **2**, 61-65.

Hasbrouck, J. and Schwartz, R. A. (1988) Liquidity and execution costs in equity markets, *Journal of Portfolio Management*, **14**, 10-17.

Hasbrouck, J. (1991a) Measuring the information content of stock trades, *Journal of Finance*, **66**, 179-207.

Hasbrouck, J. (1991b) The summary informativeness of stock trades: an econometric analysis, *Review of Financial Studies*, **4**, 571-595.

Hasbrouck, J. and Seppi, D. J. (2001) Common Factors in Prices, Order Flows and Liquidity, *Journal of Financial Economics*, **59,** 383-411.

Hausman, J. A., Lo, A. W. and McKinlay, A. C. (1992) An ordered probit analysis of transaction stock prices, *Journal of Financial Economics*, **31**, 319-379.

Huberman, G. and Stanzl, W. (2005) Optimal liquidity trading, *Review of Finance,* **9,** No. 2, 165-200.

ITG Inc. (2007) Intraday parameter distributions for world equity markets, ITG Financial Engineering, *Technical Report.*

ITG Inc. (2007) ITG Risk Models - Version 3, *ITG Inc. White Paper.*

Jain, P. and Joh, G.-H. (1988) The dependence between hourly prices and daily volume, *Journal of Financial and Quantitative Analysis*, **23**, 269-284.

Karpoff, J. M. (1988) Costly short sales and the correlation of returns with volume, *Journal of Financial Research*, **11**, 173-188.

Keim, D.B. and Madhavan, A. (1996) The upstairs market for large-block transactions: analysis and measurement of price effects, *Review of Financial Studies*, **9**, 1-36.

Kempf, A. and Korn, O. (1999) Market depth and order size, *Journal of Financial Markets*, **2**, 29-48.

Kissell, R. and Glantz, M. (2003) Optimal trading strategies: quantitative approaches for managing market impact and trading risk, *American Management Association*.

Kyle, A. S. (1985) Continuous auctions and insider trading, *Econometrica*, **53**, 1315-1336.

Lee, C., Mucklow, B. and Ready, M. (1993) Spreads, depths and the impact of earnings information: an intraday analysis, *Review of Financial Studies*, **6**, 345-374.

Lert, P. (2001) Methods of Measuring Transaction Costs, *Investment Guides*, Spring 2001.

Lo, A. and MacKinlay (1988) Stock market prices do not follow random walks: evidence from a simple specification test, *Review of Financial Studies*, **1**, 203-238.

Madhavan, A. and Smidt, S. (1991) A Bayesian model of inraday specialist pricing, *Journal of Financial Economics*, **30**, 99-134.

Madhavan, A. (2002) VWAP Strategies, *Investment Guides*, Spring 2002.

McDonald, J. B. and Newey, W. K. (1988) Partially adaptive estimation of regression models via the generalized t distribution, Econometric Theory, **4**, 428-457.

McInish, T. and Wood, R. (1992) An analysis of intraday patterns in bid/ask spreads for NYSE stocks, *Journal of Finance*, **47**, 753-764.

Munck, N. H. (2006) When Transactions Went High-Tech: A Cross-Sectional Study of Equity Trading Costs in Light of More Sophisticated Trading Systems, *Journal of Trading*, Fall 2006, 90-103.

Perold, A. (1988) The Implementation Shortfall: Paper versus Reality, *Journal of Portfolio Management*, **14**, 4-9.

 

Petersen, M. and Fialkowski, D. (1994) Posted versus effective spreads: good prices or bad quotes, *Journal of Financial Economics*, **35**, 269-292.

Ready, M. J. (1999) The specialist's discretion: stopped orders and price improvement, *Review of Financial Studies*, **12**, 1075-1112.

Rhodes-Kropf, M. (2005) Price improvement in dealership markets, *Journal of Business,* **78**, 1137-1172.

Schwartz, R.A. and Steil, B. (2002) Controlling institutional trading costs: we have met the enemy and it is us, *Journal of Portfolio Management*, **28**, 39-49.

Seppi, D. J. (1990) Equilibrium block trading and asymmetric information, *Journal of Finance*, **45**, 73-94.

Spierdijk, L., Nijman, T. E., van Soest, A. H. O. (2004) Price dynamics and trading volume: a semiparametric approach, *http://papers.ssrn.com/sol3/papers.cfm?abstract_id=512342*.

Teitelbaum, R. (2003) Know a Fund's' Cost? Look Deeper, *The New York Times Company*, February 9, 2003.

Theodossiou, P. (1998) Financial Data and the Skewed Generalized T Distribution, *Management Science*, **44**, 1650-1661.

"Transaction Costs – A Cutting-Edge Guide to Best Execution" (2001) *Investment Guides*, Spring 2001, edited by Brian R. Bruce, Institutional Investor Inc.

"Transaction Performance – The Changing Face of Trading" (2002) *Investment Guides*, Spring 2002, edited by Brian R. Bruce, Institutional Investor Inc.

Yang, J. and Jiu, B. (2006) Algorithm selection: a quantitative approach, *Algorithmic Trading II*, edited by Brian R. Bruce, Institutional Investor Inc.

Zhang, Y. (1999) Toward a theory of marginally efficient markets, *Physica A*, **269**, 30-44.

 

**TABLE 1**

This table lists the countries covered by ITG ACE/1 and ITG ACE/2, respectively. The ITG ACE/1 model is available for 42 countries and the ITG ACE/2 model is available for 21 of these countries. ITG ACE Discretionary and ITG ACE Non-Discretionary are available for ITG ACE/2 only, that is, for 21 countries.

| Country Name | ISO Country Code | ITG ACE/1 available | ITG ACE/2 available |
|---|---|---|---|
| Australia | AUS | X | X |
| Austria | AUT | X | - |
| Belgium | BEL | X | X |
| Bermuda | BMU | X | - |
| Brazil | BRA | X | - |
| Canada | CAN | X | X |
| China | CHN | X | - |
| Czech Republic | CZE | X | - |
| Denmark | DNK | X | X |
| Finland | FIN | X | X |
| France | FRA | X | X |
| Germany | DEU | X | X |
| Greece | GRC | X | - |
| Hong Kong | HKG | X | X |
| Hungary | HUN | X | - |
| India | IND | X | - |
| Indonesia | IDN | X | - |
| Ireland | IRL | X | - |
| Israel | ISR | X | - |
| Italy | ITA | X | X |
| Japan | JPN | X | X |
| Korea | KOR | X | X |
| Luxemburg | LUX | X | - |
| Malaysia | MYS | X | - |
| Mexico | MEX | X | - |
| The Netherlands | NLD | X | X |
| New Zealand | NZL | X | - |
| Norway | NOR | X | X |
| Philippines | PHL | X | - |
| Poland | POL | X | - |
| Portugal | PRT | X | X |
| Russia | RUS | X | - |
| Singapore | SGP | X | X |
| South Africa | ZAF | X | - |
| Spain | ESP | X | X |
| Sweden | SWE | X | X |
| Switzerland | CHE | X | X |
| Taiwan | TWN | X | X |
| Thailand | THA | X | - |
| Turkey | TUR | X | - |
| United Kingdom (U.K.) | GBR | X | X |
| United States (U.S.) | USA | X | X |
| **Total number** | | **42** | **21** |

 

**TABLE 2**

This table illustrates the expected trading costs, standard deviation of trading costs, and trading horizons for different values of risk aversion for ITG ACE Discretionary and ITG ACE Non-Discretionary, respectively. The underlying order is to buy a) 300,000 shares (approximately 8.5% of ADV) or b) 1,500,000 shares (approximately 42.5% of ADV) of stock YYY (a large-cap stock). The cost estimates are based on ITG ACE Optimal Strategies. Panel A reports values in cents and Panel B in basis points. ITG ACE is computed based on information as of May 1, 2007.

| | | ITG ACE Discretionary | | | ITG ACE Non-Discretionary | | |
|---|---|---|---|---|---|---|---|
| **Order Size (in shares)** | **Risk Aversion** | **Optimal Trading Horizon (in days)** | **Expected Transaction Costs (per share)** | **Standard Deviation of Transaction Costs (per share)** | **Optimal Trading Horizon (in days)** | **Expected Transaction Costs (per share)** | **Standard Deviation of Transaction Costs (per share)** |
| *Panel A. Values are in cents* | | | | | | | |
| 300,000 | 0.00 | 1 | 13.2 | 76.9 | 1 | 18.9 | 75.2 |
| 300,000 | 0.30 | 1 | 13.3 | 69.2 | 1 | 19.4 | 66.4 |
| 300,000 | 0.60 | 1 | 14.3 | 55.9 | 1 | 21.9 | 54.5 |
| 300,000 | 0.90 | 1 | 18.7 | 34.8 | 1 | 35.1 | 34.8 |
| 300,000 | 0.95 | 1 | 21.5 | 29.2 | 1 | 44.1 | 29.3 |
| 300,000 | 1.00 | 1 | 30.9 | 23.5 | 1 | 72.5 | 23.5 |
| 1,500,000 | 0.00 | 3 | 25.6 | 143.9 | 5 | 35.8 | 187.2 |
| 1,500,000 | 0.30 | 2 | 28.5 | 101.1 | 3 | 41.7 | 114.4 |
| 1,500,000 | 0.60 | 2 | 31.3 | 82.4 | 1 | 67.3 | 66.2 |
| 1,500,000 | 0.90 | 1 | 57.3 | 43.8 | 1 | 116.1 | 43.9 |
| 1,500,000 | 0.95 | 1 | 72.9 | 36.2 | 1 | 156.1 | 36.2 |
| 1,500,000 | 1.00 | 1 | 123.9 | 26.5 | 1 | 289.1 | 26.5 |
| *Panel B. Values are in basis points* | | | | | | | |
| 300,000 | 0.00 | 1 | 14.1 | 82.5 | 1 | 20.3 | 80.6 |
| 300,000 | 0.30 | 1 | 14.3 | 74.2 | 1 | 20.8 | 71.3 |
| 300,000 | 0.60 | 1 | 15.3 | 60.0 | 1 | 23.5 | 58.4 |
| 300,000 | 0.90 | 1 | 20.0 | 37.3 | 1 | 37.7 | 37.3 |
| 300,000 | 0.95 | 1 | 23.0 | 31.4 | 1 | 47.3 | 31.5 |
| 300,000 | 1.00 | 1 | 33.2 | 25.2 | 1 | 77.8 | 25.2 |
| 1,500,000 | 0.00 | 3 | 27.5 | 154.4 | 5 | 38.4 | 200.8 |
| 1,500,000 | 0.30 | 2 | 30.5 | 108.5 | 3 | 44.7 | 122.7 |
| 1,500,000 | 0.60 | 2 | 33.6 | 88.4 | 1 | 72.1 | 71.0 |
| 1,500,000 | 0.90 | 1 | 61.5 | 47.0 | 1 | 124.6 | 47.0 |
| 1,500,000 | 0.95 | 1 | 78.2 | 38.8 | 1 | 167.5 | 38.9 |
| 1,500,000 | 1.00 | 1 | 132.9 | 28.4 | 1 | 310.1 | 28.4 |

 

**TABLE 3**

This table reports descriptive statistics of the data for the calibration/testing of the ITG ACE model for some of the markets in the ITG ACE universe. The statistics are based on the time period from January 2004 to December 2006. Reported are the number of executions, the number of clusters (or order decisions), the volume of the executions in local currency, the number of stocks executions are recorded for, and the number of clients in ITG's Peer Group Database. The countries are sorted by decreasing number of executions.

| Country Name | ISO Country Code | Number of Executions (in million) | Number of Clusters (Order Decisions) | Local Currency Volume (in billion) | Local Currency | Number of Stocks | Number of Clients |
|---|---|---|---|---|---|---|---|
| United States (U.S.) | USA | 42.300 | 4,592,255 | 6,402.0 | USD | 7,405 | 89 |
| United Kingdom (U.K.) | GBR | 2.110 | 290,093 | 415.0 | GBP | 1,871 | 45 |
| Japan | JPN | 1.530 | 201,130 | 419.0 | JPY | 1,794 | 48 |
| France | FRA | 0.800 | 127,390 | 233.9 | EUR | 477 | 49 |
| Germany | DEU | 0.570 | 98,279 | 176.9 | EUR | 377 | 49 |
| Switzerland | CHE | 0.470 | 69,302 | 210.3 | CHF | 223 | 48 |
| The Netherlands | NLD | 0.390 | 58,530 | 103.2 | EUR | 119 | 48 |
| Italy | ITA | 0.270 | 54,069 | 76.0 | EUR | 190 | 46 |
| Canada | CAN | 0.270 | 46,535 | 81.3 | CAD | 739 | 51 |
| Spain | ESP | 0.220 | 42,264 | 63.5 | EUR | 98 | 47 |
| Australia | AUS | 0.190 | 38,600 | 62.4 | AUD | 399 | 46 |
| Sweden | SWE | 0.200 | 27,415 | 444.7 | SEK | 168 | 46 |
| Finland | FIN | 0.110 | 19,980 | 30.1 | EUR | 76 | 40 |
| Norway | NOR | 0.068 | 13,435 | 156.1 | NOK | 124 | 42 |
| Belgium | BEL | 0.063 | 13,228 | 14.4 | EUR | 74 | 44 |
| Denmark | DNK | 0.051 | 9,007 | 96.0 | DKK | 80 | 45 |
| Portugal | PRT | 0.022 | 4,284 | 5.3 | EUR | 22 | 35 |

## TABLE 4

This table reports results for tests of the ITG ACE/2 Discretionary model for the U.S. The test results are based on the sign and the Wilcoxon signed-rank tests. The tests are based on the time period from January 2005 to December 2006. The stocks are divided into Listed (NYSE and AMEX) stocks and OTC stocks. We then consider four groups for Listed/OTC stocks: large-cap, mid-cap, small-cap, and all stocks. For each stock group, orders are grouped in terms of order size according to three subdivisions:

1. Uniform Subdivision 1: 0-4%, 4-8%, 8-12%, 12-16%, 16-20%, 20-24% of average daily volume (ADV)
2. Uniform Subdivision 2: 0-15%, 15-30%, 30-45%, 45-60%, 60-75%, 75-90% of ADV
3. ITG TCA[16] Subdivision: 0-1%, 1-5%, 5-10%, 10-25%, 25-50%, > 50% of ADV

Each of the groups of a subdivision is then further subdivided into ten equally-sized intervals (e.g., Group 0-15% is subdivided into 0-1.5%, 1.5-3%, …, 13.5-15% of ADV). For each of these intervals, the average empirical and the average ITG ACE transaction costs and the difference between the two averages are computed. We then calculate the sign and the Wilcoxon signed-rank test statistics of the difference based on the ten observations[17] (since we have ten intervals for each group of a subdivision). The null hypothesis that the average empirical and the average ITG ACE transaction costs are equal for a specific group of a subdivision is rejected if the respective test statistic is significant at the 5% level. Note, there are 96 possible rejections of the null hypothesis: 2 (exchanges) x 4 (stock groups) x 6 (subdivision groups) x 2 (tests).

| **Subdivision** | **Number of Rejections** | **Total Possible Rejections** | **Proportion Rejected (in %)** |
|---|---|---|---|
| *Overall (Over- and Underestimation)* | | | |
| Uniform Subdivision 1 – EW | 3 | 96 | 3.13 |
| Uniform Subdivision 1 – VW | 4 | 96 | 4.17 |
| Uniform Subdivision 2 – EW | 6 | 96 | 6.25 |
| Uniform Subdivision 2 – VW | 8 | 96 | 8.33 |
| ITG TCA Subdivision – EW | 9 | 96 | 9.38 |
| ITG TCA Subdivision – VW | 8 | 96 | 8.33 |
| *Overestimation (average ITG ACE costs are higher than the average empirical costs)* | | | |
| Uniform Subdivision 1 – EW | 2 | 96 | 2.08 |
| Uniform Subdivision 1 – VW | 3 | 96 | 3.13 |
| Uniform Subdivision 2 – EW | 4 | 96 | 4.17 |
| Uniform Subdivision 2 – VW | 4 | 96 | 4.17 |
| ITG TCA Subdivision – EW | 5 | 96 | 5.21 |
| ITG TCA Subdivision – VW | 3 | 96 | 3.13 |
| *Underestimation (average ITG ACE costs are lower than the average empirical costs)* | | | |
| Uniform Subdivision 1 – EW | 1 | 96 | 1.04 |
| Uniform Subdivision 1 – VW | 1 | 96 | 1.04 |
| Uniform Subdivision 2 – EW | 2 | 96 | 2.08 |
| Uniform Subdivision 2 – VW | 4 | 96 | 4.17 |
| ITG TCA Subdivision – EW | 4 | 96 | 4.17 |
| ITG TCA Subdivision – VW | 5 | 96 | 5.21 |

Note:
- EW denotes that orders are weighted equally in computing the test statistics.
- VW denotes that orders are weighted by the total order value in computing the test statistics.
- The null hypothesis of the sign test is that the median of the cost differences is zero.
- The null hypothesis of the Wilcoxon signed-rank test is that the cost differences are distributed symmetrically around zero.

[16] ITG TCA stands for ITG Transaction Cost Analysis and is ITG's tool for post-trade analysis.
[17] The number of observations can be less than ten, if there are no orders in a specific interval.

 

## TABLE 5

This table reports results for tests of the ITG ACE/2 Discretionary model for Canada. The test results are based on the sign and the Wilcoxon signed-rank tests. The tests are based on the time period from January 2005 to December 2006. We consider four groups of stocks: large-cap, mid-cap, small-cap, and all stocks. For each stock group, orders are grouped in terms of order size according to three subdivisions:

1. Uniform Subdivision 1: 0-4%, 4-8%, 8-12%, 12-16%, 16-20%, 20-24% of average daily volume (ADV)
2. Uniform Subdivision 2: 0-15%, 15-30%, 30-45%, 45-60%, 60-75%, 75-90% of ADV
3. ITG TCA[18] Subdivision: 0-1%, 1-5%, 5-10%, 10-25%, 25-50%, > 50% of ADV

Each of the groups of a subdivision is then further subdivided into ten equally-sized intervals (e.g., Group 0-15% is subdivided into 0-1.5%, 1.5-3%, …, 13.5-15% of ADV). For each of these intervals, the average empirical and the average ITG ACE transaction costs and the difference between the two averages are computed. We then calculate the sign and the Wilcoxon signed-rank test statistics of the difference based on the ten observations[19] (since we have ten intervals for each group of a subdivision). The null hypothesis that the average empirical and the average ITG ACE transaction costs are equal for a specific group of a subdivision is rejected if the respective test statistic is significant at the 5% level. Note, there are 48 possible rejections of the null hypothesis: 4 (stock groups) x 6 (subdivision groups) x 2 (tests).

| Subdivision | Number of Rejections | Total Possible Rejections | Proportion Rejected (in %) |
|---|---|---|---|
| *Overall (Over- and Underestimation)* | | | |
| Uniform Subdivision 1 – EW | 5 | 48 | 10.42 |
| Uniform Subdivision 1 – VW | 0 | 48 | 0.00 |
| Uniform Subdivision 2 – EW | 6 | 48 | 12.50 |
| Uniform Subdivision 2 – VW | 3 | 48 | 6.25 |
| ITG TCA Subdivision – EW | 11 | 48 | 22.92 |
| ITG TCA Subdivision – VW | 6 | 48 | 12.50 |
| *Overestimation (average ITG ACE costs are higher than the average empirical costs)* | | | |
| Uniform Subdivision 1 – EW | 5 | 48 | 10.42 |
| Uniform Subdivision 1 – VW | 0 | 48 | 0.00 |
| Uniform Subdivision 2 – EW | 6 | 48 | 12.50 |
| Uniform Subdivision 2 – VW | 3 | 48 | 6.25 |
| ITG TCA Subdivision – EW | 9 | 48 | 18.75 |
| ITG TCA Subdivision – VW | 4 | 48 | 8.33 |
| *Underestimation (average ITG ACE costs are lower than the average empirical costs)* | | | |
| Uniform Subdivision 1 – EW | 0 | 48 | 0.00 |
| Uniform Subdivision 1 – VW | 0 | 48 | 0.00 |
| Uniform Subdivision 2 – EW | 0 | 48 | 0.00 |
| Uniform Subdivision 2 – VW | 0 | 48 | 0.00 |
| ITG TCA Subdivision – EW | 2 | 48 | 4.17 |
| ITG TCA Subdivision – VW | 2 | 48 | 4.17 |

Note:
- EW denotes that orders are weighted equally in computing the test statistics.
- VW denotes that orders are weighted by the total order value in computing the test statistics.
- The null hypothesis of the sign test is that the median of the cost differences is zero.
- The null hypothesis of the Wilcoxon signed-rank test is that the cost differences are distributed symmetrically around zero.

[18] ITG TCA stands for ITG Transaction Cost Analysis and is ITG's tool for post-trade analysis.
[19] The number of observations can be less than ten, if there are no orders in a specific interval.

## TABLE 6

This table reports results for tests of the ITG ACE/2 Discretionary model for the United Kingdom (U.K.). The test results are based on the sign and the Wilcoxon signed-rank tests. The tests are based on the time period from January 2005 to December 2006. The stocks are divided into SETS and SEAQ stocks. We then consider four groups for SETS/SEAQ: large-cap, mid-cap, small-cap, and all stocks. For each stock group, orders are grouped in terms of order size according to three subdivisions:

1. Uniform Subdivision 1: 0-4%, 4-8%, 8-12%, 12-16%, 16-20%, 20-24% of average daily volume (ADV)
2. Uniform Subdivision 2: 0-15%, 15-30%, 30-45%, 45-60%, 60-75%, 75-90% of ADV
3. ITG TCA[20] Subdivision: 0-1%, 1-5%, 5-10%, 10-25%, 25-50%, > 50% of ADV

Each of the groups of a subdivision is then further subdivided into ten equally-sized intervals (e.g., Group 0-15% is subdivided into 0-1.5%, 1.5-3%, …, 13.5-15% of ADV). For each of these intervals, the average empirical and the average ITG ACE transaction costs and the difference between the two averages are computed. We then calculate the sign and the Wilcoxon signed-rank test statistics of the difference based on the ten observations[21] (since we have ten intervals for each group of a subdivision). The null hypothesis that the average empirical and the average ITG ACE transaction costs are equal for a specific group of a subdivision is rejected if the respective test statistic is significant at the 5% level. Note, there are 96 possible rejections of the null hypothesis: 2 (exchanges) x 4 (stock groups) x 6 (subdivision groups) x 2 (tests).

| Subdivision | Number of Rejections | Total Possible Rejections | Proportion Rejected (in %) |
|---|---|---|---|
| *Overall (Over- and Underestimation)* | | | |
| Uniform Subdivision 1 – EW | 5 | 96 | 5.21 |
| Uniform Subdivision 1 – VW | 0 | 96 | 0.00 |
| Uniform Subdivision 2 – EW | 6 | 96 | 6.25 |
| Uniform Subdivision 2 – VW | 17 | 96 | 17.71 |
| ITG TCA Subdivision – EW | 13 | 96 | 13.54 |
| ITG TCA Subdivision – VW | 11 | 96 | 11.46 |
| *Overestimation (average ITG ACE costs are higher than the average empirical costs)* | | | |
| Uniform Subdivision 1 – EW | 3 | 96 | 3.13 |
| Uniform Subdivision 1 – VW | 0 | 96 | 0.00 |
| Uniform Subdivision 2 – EW | 0 | 96 | 0.00 |
| Uniform Subdivision 2 – VW | 2 | 96 | 2.08 |
| ITG TCA Subdivision – EW | 6 | 96 | 6.25 |
| ITG TCA Subdivision – VW | 3 | 96 | 3.13 |
| *Underestimation (average ITG ACE costs are lower than the average empirical costs)* | | | |
| Uniform Subdivision 1 – EW | 2 | 96 | 2.08 |
| Uniform Subdivision 1 – VW | 0 | 96 | 0.00 |
| Uniform Subdivision 2 – EW | 6 | 96 | 6.25 |
| Uniform Subdivision 2 – VW | 15 | 96 | 15.63 |
| ITG TCA Subdivision – EW | 7 | 96 | 7.29 |
| ITG TCA Subdivision – VW | 8 | 96 | 8.33 |

Note:
- EW denotes that orders are weighted equally in computing the test statistics.
- VW denotes that orders are weighted by the total order value in computing the test statistics.
- The null hypothesis of the sign test is that the median of the cost differences is zero.
- The null hypothesis of the Wilcoxon signed-rank test is that the cost differences are distributed symmetrically around zero.

[20] ITG TCA stands for ITG Transaction Cost Analysis and is ITG's tool for post-trade analysis.
[21] The number of observations can be less than ten, if there are no orders in a specific interval.

 

**TABLE 7**

This table reports results for tests of the ITG ACE/2 Discretionary model for Japan. The test results are based on the sign and the Wilcoxon signed-rank tests. The tests are based on the time period from January 2005 to December 2006. The stocks are divided into Tokyo Stock Exchange (TSE) stocks and non-TSE stocks. We then consider four groups for TSE/non-TSE: large-cap, mid-cap, small-cap, and all stocks. For each stock group, orders are grouped in terms of order size according to three subdivisions:

4. Uniform Subdivision 1: 0-4%, 4-8%, 8-12%, 12-16%, 16-20%, 20-24% of average daily volume (ADV)
5. Uniform Subdivision 2: 0-15%, 15-30%, 30-45%, 45-60%, 60-75%, 75-90% of ADV
6. ITG TCA[22] Subdivision: 0-1%, 1-5%, 5-10%, 10-25%, 25-50%, > 50% of ADV

Each of the groups of a subdivision is then further subdivided into ten equally-sized intervals (e.g., Group 0-15% is subdivided into 0-1.5%, 1.5-3%, …, 13.5-15% of ADV). For each of these intervals, the average empirical and the average ITG ACE transaction costs and the difference between the two averages are computed. We then calculate the sign and the Wilcoxon signed-rank test statistics of the difference based on the ten observations[23] (since we have ten intervals for each group of a subdivision). The null hypothesis that the average empirical and the average ITG ACE transaction costs are equal for a specific group of a subdivision is rejected if the respective test statistic is significant at the 5% level. Note, there are 96 possible rejections of the null hypothesis: 2 (exchanges) x 4 (stock groups) x 6 (subdivision groups) x 2 (tests).

| Subdivision | Number of Rejections | Total Possible Rejections | Proportion Rejected (in %) |
|---|---|---|---|
| *Overall (Over- and Underestimation)* | | | |
| Uniform Subdivision 1 – EW | 2 | 96 | 2.08 |
| Uniform Subdivision 1 – VW | 2 | 96 | 2.08 |
| Uniform Subdivision 2 – EW | 13 | 96 | 13.54 |
| Uniform Subdivision 2 – VW | 15 | 96 | 5.21 |
| ITG TCA Subdivision – EW | 10 | 96 | 10.43 |
| ITG TCA Subdivision – VW | 8 | 96 | 8.35 |
| *Overestimation (average ITG ACE costs are higher than the average empirical costs)* | | | |
| Uniform Subdivision 1 – EW | 2 | 96 | 2.08 |
| Uniform Subdivision 1 – VW | 2 | 96 | 2.08 |
| Uniform Subdivision 2 – EW | 12 | 96 | 12.50 |
| Uniform Subdivision 2 – VW | 3 | 96 | 3.13 |
| ITG TCA Subdivision – EW | 7 | 96 | 7.29 |
| ITG TCA Subdivision – VW | 2 | 96 | 2.08 |
| *Underestimation (average ITG ACE costs are lower than the average empirical costs)* | | | |
| Uniform Subdivision 1 – EW | 0 | 96 | 0.00 |
| Uniform Subdivision 1 – VW | 0 | 96 | 0.00 |
| Uniform Subdivision 2 – EW | 1 | 96 | 1.04 |
| Uniform Subdivision 2 – VW | 2 | 96 | 2.08 |
| ITG TCA Subdivision – EW | 3 | 96 | 3.13 |
| ITG TCA Subdivision – VW | 6 | 96 | 6.27 |

Note:
- EW denotes that orders are weighted equally in computing the test statistics.
- VW denotes that orders are weighted by the total order value in computing the test statistics.
- The null hypothesis of the sign test is that the median of the cost differences is zero.
- The null hypothesis of the Wilcoxon signed-rank test is that the cost differences are distributed symmetrically around zero.

[22] ITG TCA™ stands for ITG Transaction Cost Analysis and is ITG's tool for post-trade analysis.
[23] The number of observations can be less than ten, if there are no orders in a specific interval.

## TABLE 8

This table reports results for tests of the ITG ACE/2 Discretionary model for Euronext, consisting of the exchanges in France, Belgium, The Netherlands, and Portugal. The test results are based on the sign and the Wilcoxon signed-rank tests. The tests are based on the time period from January 2005 to December 2006. The stocks are divided by country. We then consider four groups for each country: large-cap, mid-cap, small-cap, and all stocks. For each stock group, orders are grouped in terms of order size according to three subdivisions:

7. Uniform Subdivision 1: 0-4%, 4-8%, 8-12%, 12-16%, 16-20%, 20-24% of average daily volume (ADV)
8. Uniform Subdivision 2: 0-15%, 15-30%, 30-45%, 45-60%, 60-75%, 75-90% of ADV
9. ITG TCA[24] Subdivision: 0-1%, 1-5%, 5-10%, 10-25%, 25-50%, > 50% of ADV

Each of the groups of a subdivision is then further subdivided into ten equally-sized intervals (e.g., Group 0-15% is subdivided into 0-1.5%, 1.5-3%, …, 13.5-15% of ADV). For each of these intervals, the average empirical and the average ITG ACE transaction costs and the difference between the two averages are computed. We then calculate the sign and the Wilcoxon signed-rank test statistics of the difference based on the ten observations[25] (since we have ten intervals for each group of a subdivision). The null hypothesis that the average empirical and the average ITG ACE transaction costs are equal for a specific group of a subdivision is rejected if the respective test statistic is significant at the 5% level. Note, there are192 possible rejections of the null hypothesis: 4 (countries) x 4 (stock groups) x 6 (subdivision groups) x 2 (tests).

| Subdivision | Number of Rejections | Total Possible Rejections | Proportion Rejected (in %) |
|---|---|---|---|
| *Overall (Over- and Underestimation)* | | | |
| Uniform Subdivision 1 – EW | 7 | 192 | 3.65 |
| Uniform Subdivision 1 – VW | 3 | 192 | 1.56 |
| Uniform Subdivision 2 – EW | 16 | 192 | 8.33 |
| Uniform Subdivision 2 – VW | 5 | 192 | 2.60 |
| ITG TCA Subdivision – EW | 23 | 192 | 11.98 |
| ITG TCA Subdivision – VW | 12 | 192 | 6.25 |
| *Overestimation (average ITG ACE costs are higher than the average empirical costs)* | | | |
| Uniform Subdivision 1 – EW | 4 | 192 | 2.08 |
| Uniform Subdivision 1 – VW | 1 | 192 | 0.52 |
| Uniform Subdivision 2 – EW | 14 | 192 | 7.29 |
| Uniform Subdivision 2 – VW | 4 | 192 | 2.08 |
| ITG TCA Subdivision – EW | 18 | 192 | 9.38 |
| ITG TCA Subdivision – VW | 10 | 192 | 5.21 |
| *Underestimation (average ITG ACE costs are lower than the average empirical costs)* | | | |
| Uniform Subdivision 1 – EW | 3 | 192 | 1.56 |
| Uniform Subdivision 1 – VW | 2 | 192 | 1.04 |
| Uniform Subdivision 2 – EW | 2 | 192 | 1.04 |
| Uniform Subdivision 2 – VW | 1 | 192 | 0.52 |
| ITG TCA Subdivision – EW | 5 | 192 | 2.60 |
| ITG TCA Subdivision – VW | 2 | 192 | 1.04 |

Note:
- EW denotes that orders are weighted equally in computing the test statistics.
- VW denotes that orders are weighted by the total order value in computing the test statistics.
- The null hypothesis of the sign test is that the median of the cost differences is zero.
- The null hypothesis of the Wilcoxon signed-rank test is that the cost differences are distributed symmetrically around zero.

[24] ITG TCA stands for ITG Transaction Cost Analysis and is ITG's tool for post-trade analysis.
[25] The number of observations can be less than ten, if there are no orders in a specific interval.

**FIGURE 1**


This figure illustrates the concept behind the ITG ACE price impact model for a sell trade. The execution price of the stock is lower than the pre-trade price as the law of supply and demand suggests. The larger the size of the trade, the more likely the sale price will be lower. The difference between pre-trade market price and execution price consists of two parts – permanent and temporary price impact. While the temporary price impact only affects the price of the trade itself, the permanent price impact has a persistent effect on the market price.

 

**FIGURE 2**

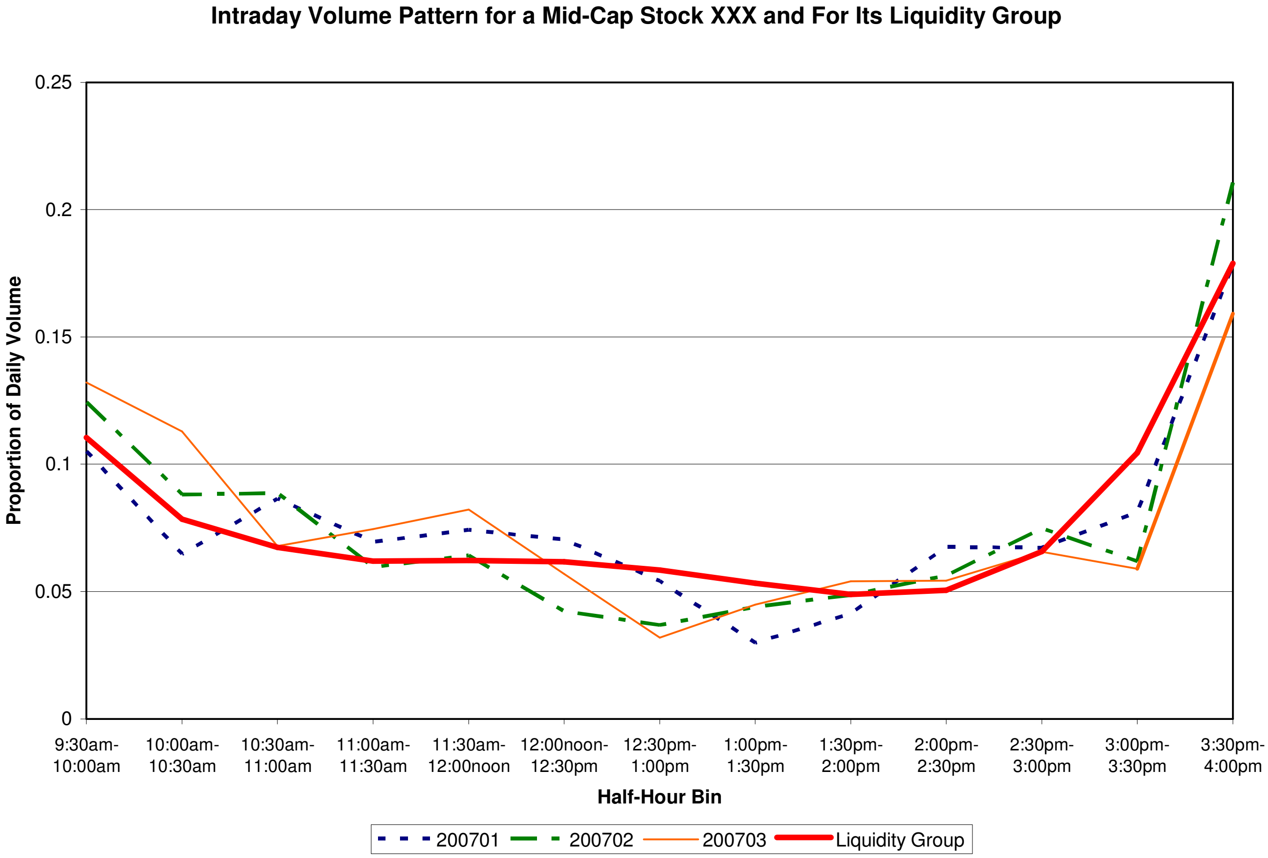


This figure shows the intraday volume pattern for stock XXX for the months January, February, and March of 2007. The stock is a randomly chosen mid-cap stock and is a relatively illiquid stock; its market capitalization is about $390 million, and the median daily share volume is about 50,000 shares. The distributions show some fluctuations, especially at the beginning and at the end of the trading day. The bold line represents the smoothed average intraday volume distributions for all stocks which belong to the same market and liquidity group as stock XXX. The average was taken over the three-month period from January to March, 2007.

**FIGURE 3**

**Intraday Bid-Ask Spread Pattern for a Mid-Cap Stock XXX and For Its Liquidity Group**

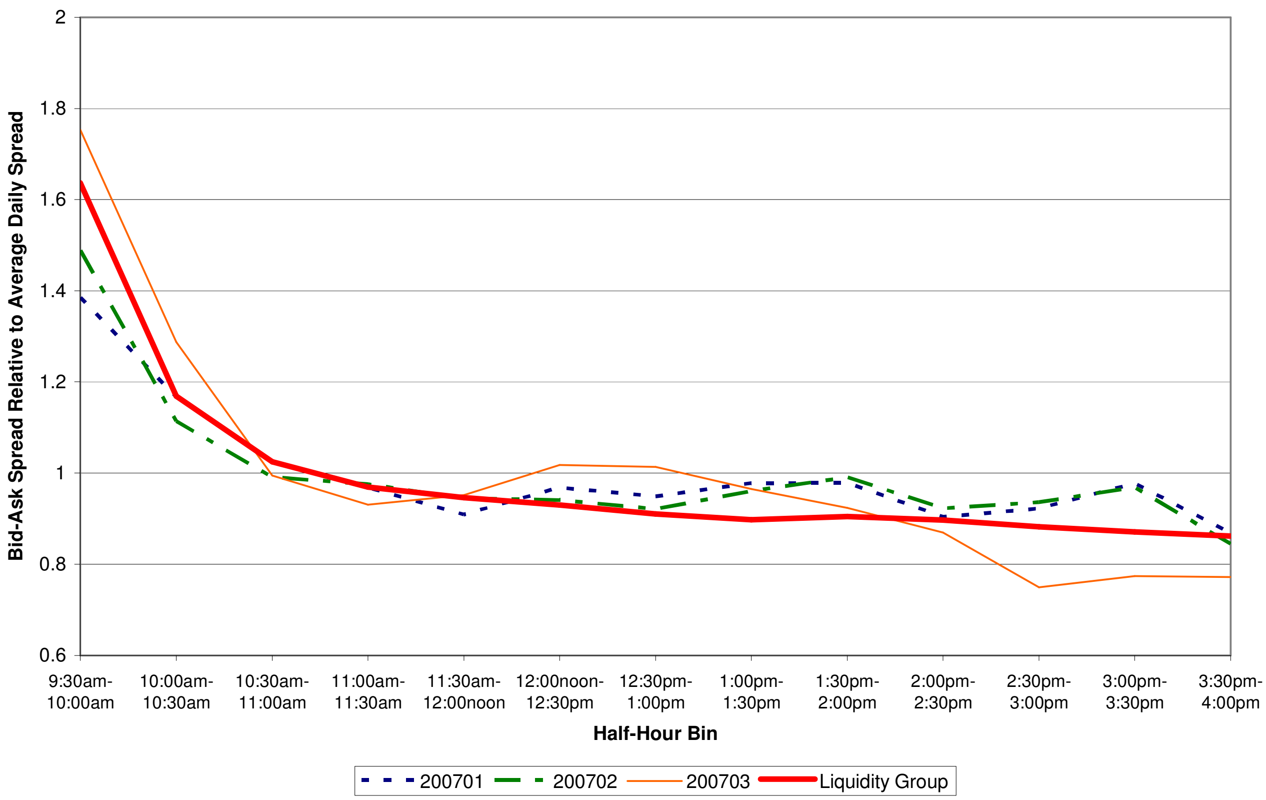


This figure shows the intraday bid-ask spread pattern for stock XXX for the months January, February, and March of 2007. The stock is a randomly chosen mid-cap stock and is a relatively illiquid stock; its market capitalization is about $390 million and the median daily share volume is about 50,000 shares. The distributions show some fluctuations, especially at the beginning and at the end of the trading day. The bold line represents the smoothed average intraday bid-ask spread distribution for all stocks which belong to the same market and liquidity group as stock XXX. The average was taken over the three-month period from January to March, 2007.

**FIGURE 4**

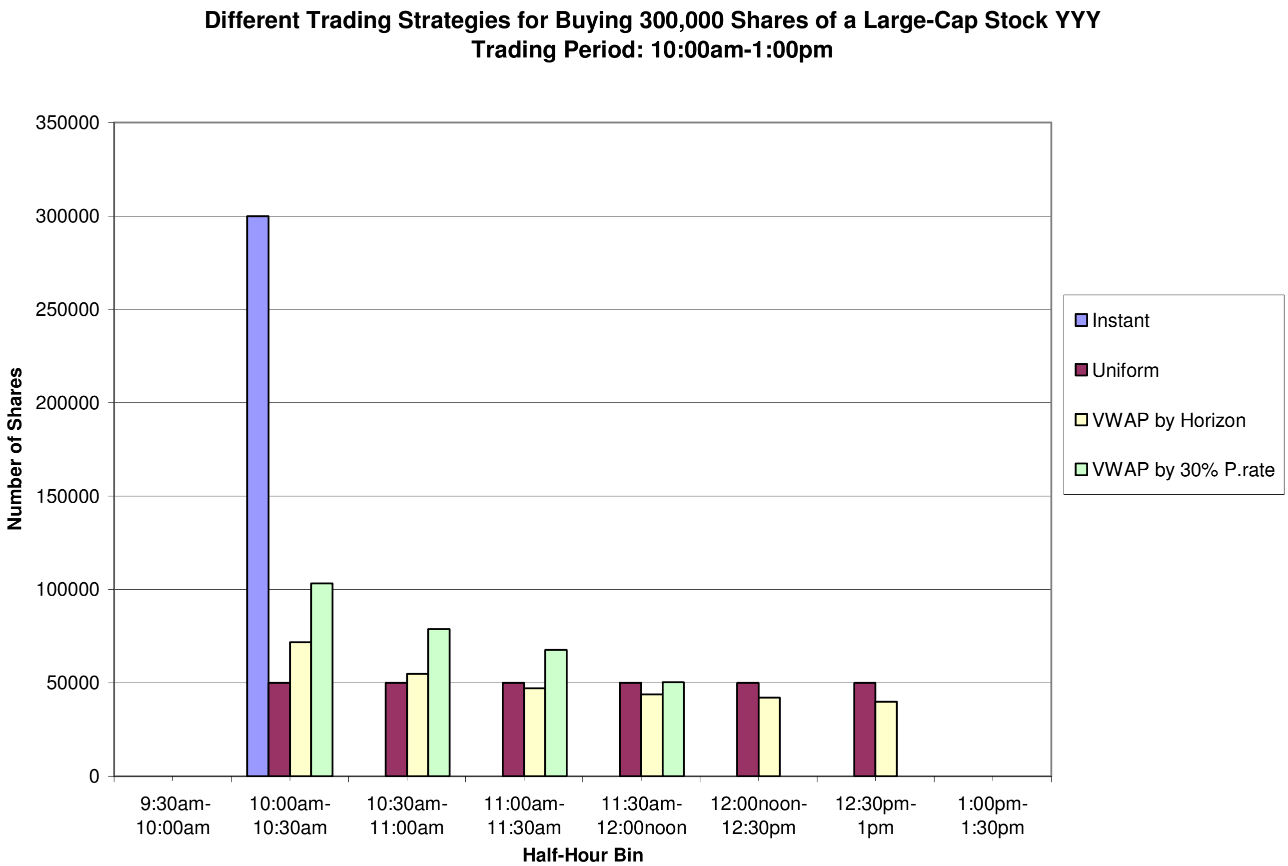


This figure shows different types of trading strategies for buying 300,000 shares (approximately 8.5% of ADV) of stock YYY between 10:00 a.m. and 1 p.m. The stock is a randomly chosen large-cap stock and is a relatively liquid stock; its market capitalization is over $70 billion and the median daily share volume is 3.5 million shares. The instant strategy places all the shares in the first trading bin (bin 2, i.e. 10:00 a.m. - 10:30 a.m.). The uniform strategy assumes the same number of shares to be executed for each bin within the trading period. The VWAP strategies by horizon and by 30% participation rate match the intraday volume pattern of the stock. As the intraday volume suggests, more shares are executed in the early morning.

## FIGURE 5

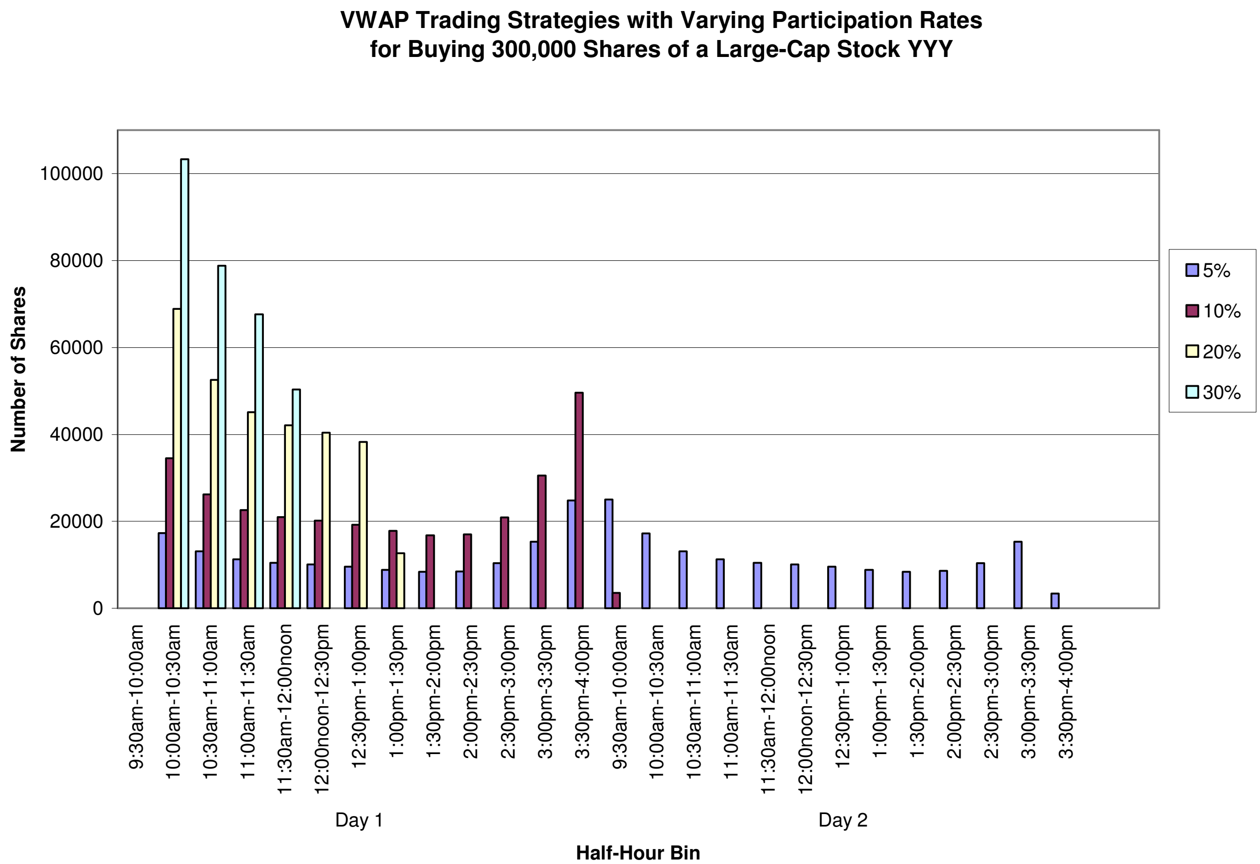


This figure shows VWAP trading strategies with varying participation rates (5%, 10%, 20%, and 30%) for buying 300,000 shares (approximately 8.5% of ADV) of stock YYY. The stock is a randomly chosen large-cap stock and is a relatively liquid stock; its market capitalization is over $70 billion and the median daily share volume is 3.5 million shares. In contrast to a VWAP trading strategy by horizon, the trade horizon is not fixed but rather depends on the participation rate. The lower the participation rate, the longer it takes to fill the order.

 

**FIGURE 6**

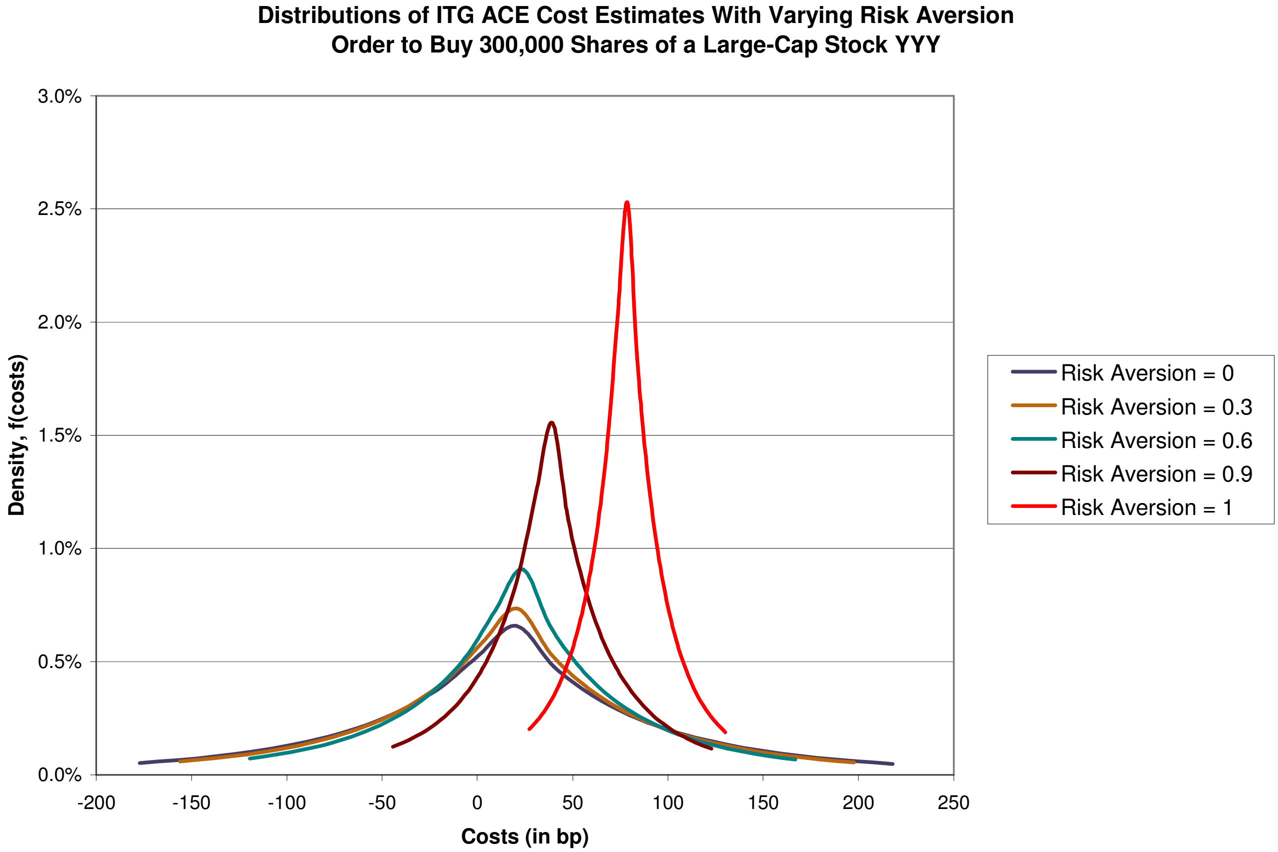


This figure illustrates the distributions of ITG ACE/2 Non-Discretionary transaction cost estimates based on different values of risk aversion (0, 0.3, 0.6, 0.9, and 1) for an order to buy 300,000 shares (approximately 8.5% of ADV) of stock YYY. The stock is a randomly chosen large-cap stock and is a relatively liquid stock; its market capitalization is over $70 billion and the median daily share volume is 3.5 million shares. The distributions are based on ITG ACE Optimal Strategies with a one-day trading horizon. The plot suggests that the choice of a greater risk aversion provides higher expected costs, but lower standard deviation of costs and, thus, potentially less opportunity costs.

**FIGURE 7**

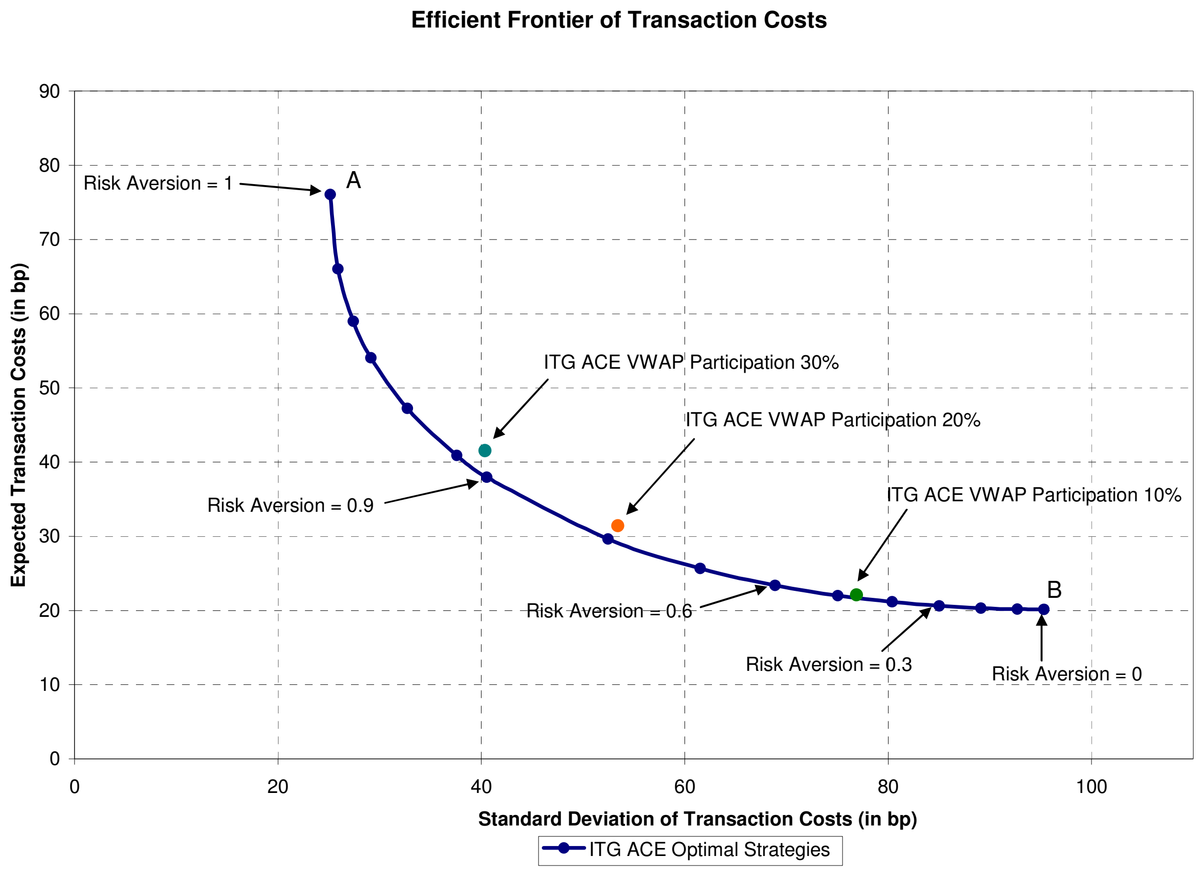


This figure graphically displays ITG ACE Optimal Strategies for different risk aversions for an order to buy 300,000 shares of a large-cap stock YYY using ITG ACE/2 Non-Discretionary. The 300,000 share order corresponds to about 8.5% of ADV. Obviously, there are many choices of optimal strategies between the two extremes of minimizing expected transaction costs (Point B) and minimizing the standard deviation of transaction costs (Point A). Each point on the efficient frontier corresponds to a specific risk aversion. The graph highlights selected risk aversion values. As you move from left to right, you can see that you can incrementally reduce the expected transaction costs of a trading strategy (relative to the most expensive) by assuming more risk. Somewhere along this "efficient frontier" of transaction costs is a strategy that, beyond which, you begin to accumulate more risk than the reduction in expected transaction costs is worth. This would be a desirable choice of risk aversion. For comparison, trading strategies other than ITG ACE Optimal Strategies are also included. As expected, theses alternative trading strategies do not lie on the efficient frontier as they are not optimal: There are trading strategies with lower expected transaction costs with the same standard deviation of transaction costs, or there are trading strategies with the same expected transaction costs, but lower standard deviation of transaction costs. Note, for all strategies the trading horizon was restricted to one trading day (with potential start in the first bin). For the VWAP By Participation Strategies, the order size is sufficiently small to ensure that the trading horizon is less than one trading day.

**FIGURE 8**

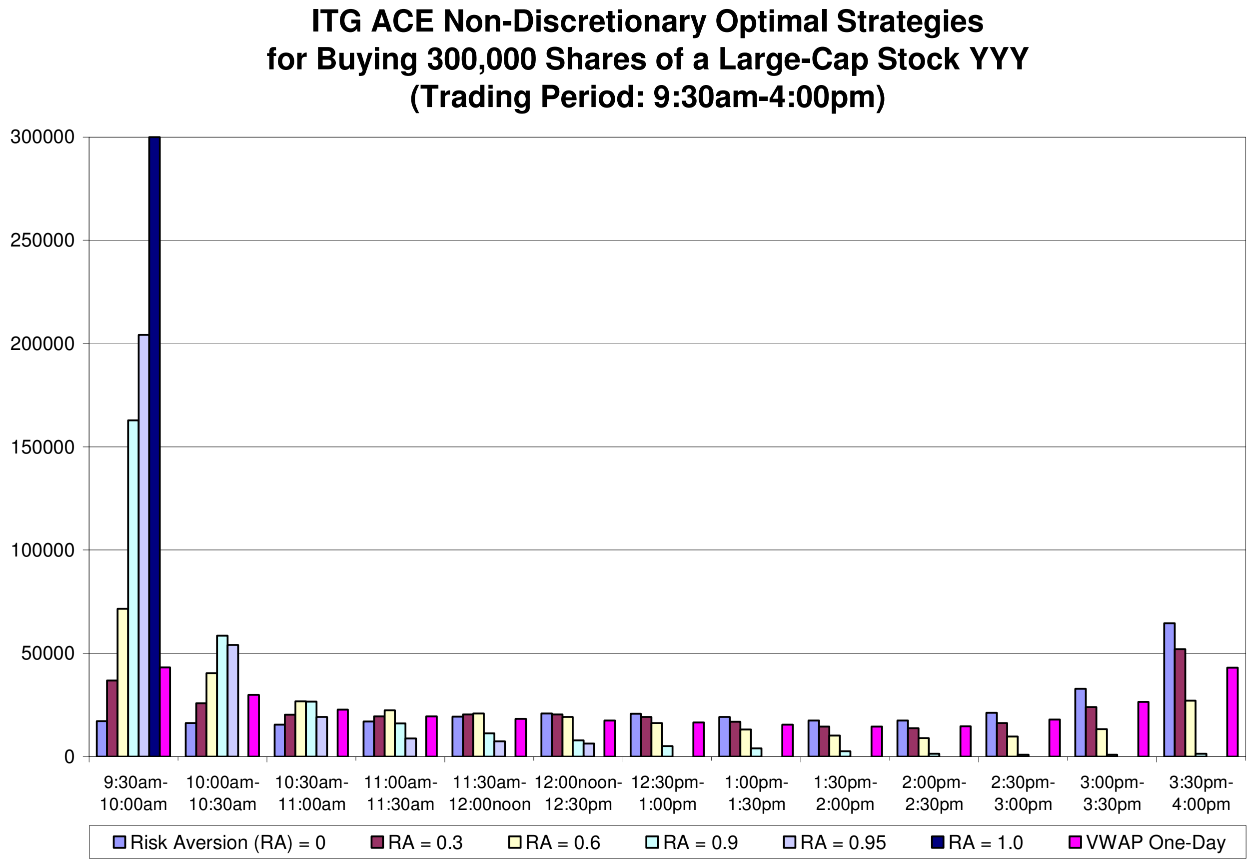


This figure shows the ITG ACE Optimal Strategies of ITG ACE/2 Non-Discretionary for buying 300,000 shares (approximately 8.5% of ADV) of a large-cap stock YYY obtained by using values of risk aversion of 0, 0.3, 0.6, 0.9, 0.95 and 1, and a one-day trading horizon. Also shown is a VWAP Strategy by Horizon with a one-day trading horizon. The ITG ACE Optimal Strategy for larger risk aversion parameters always suggests to trade more aggressively at the beginning of the trading horizon to minimize opportunity costs.

 

**FIGURE 9**

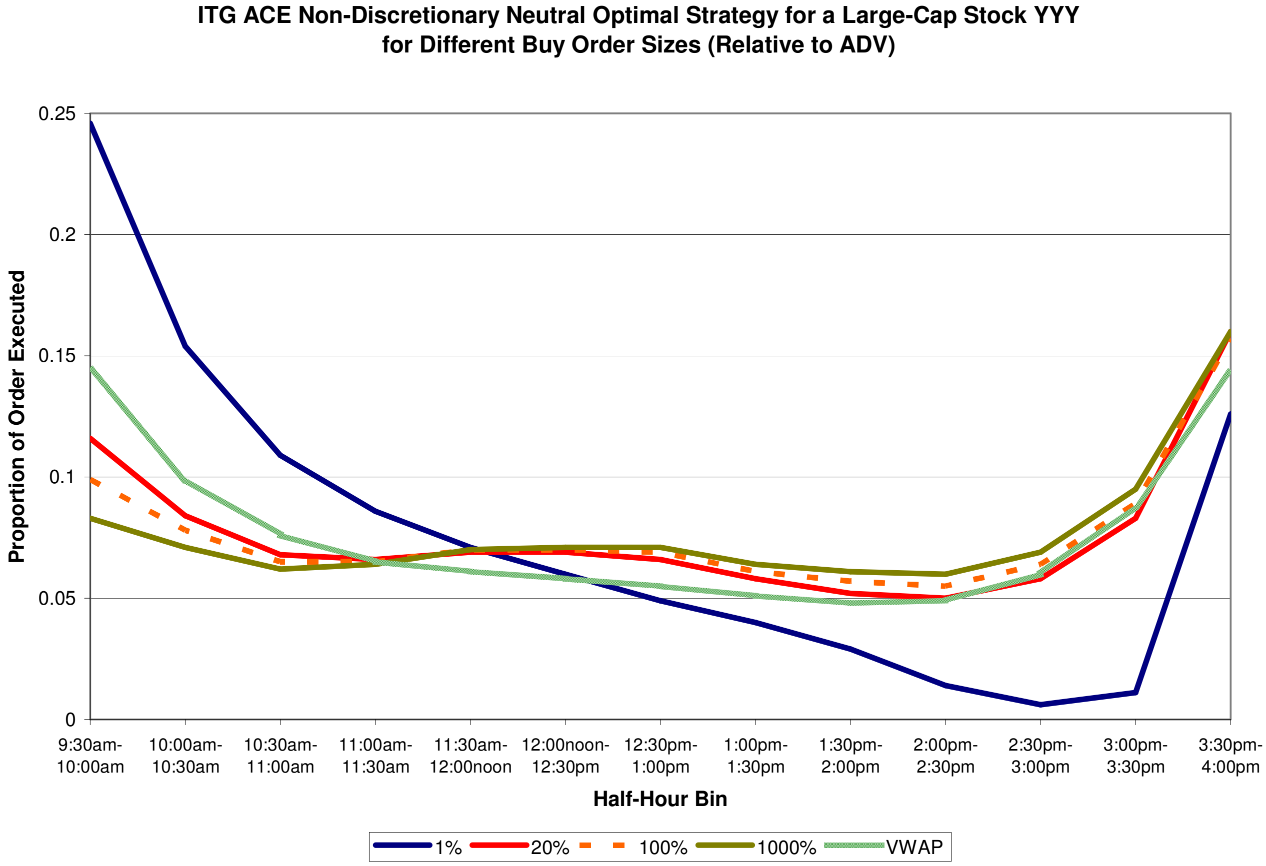


This figure illustrates different ITG ACE Optimal Strategy trading distributions for risk aversion 0.3 (ITG ACE/2 Non-Discretionary Neutral) and fixed one-day horizon for a large-cap stock YYY and different order sizes (15, 20%, 100%, and 1000% of ADV). The figure also shows the trading distribution for a one-day VWAP trading strategy. The chart shows that risk aversion 0.3 yields ITG ACE Optimal Strategies that are close to a VWAP trading strategy. Moreover, the ITG ACE Optimal Strategy becomes more and more back-loaded with increasing order size due to market impact costs.

**FIGURE 10**

**Empirical and Theoretical ITG ACE Permanent Price Impact Functions in Bin 1 (9:30am-10:00am) for the Most Liquid 20 Listed Stocks**

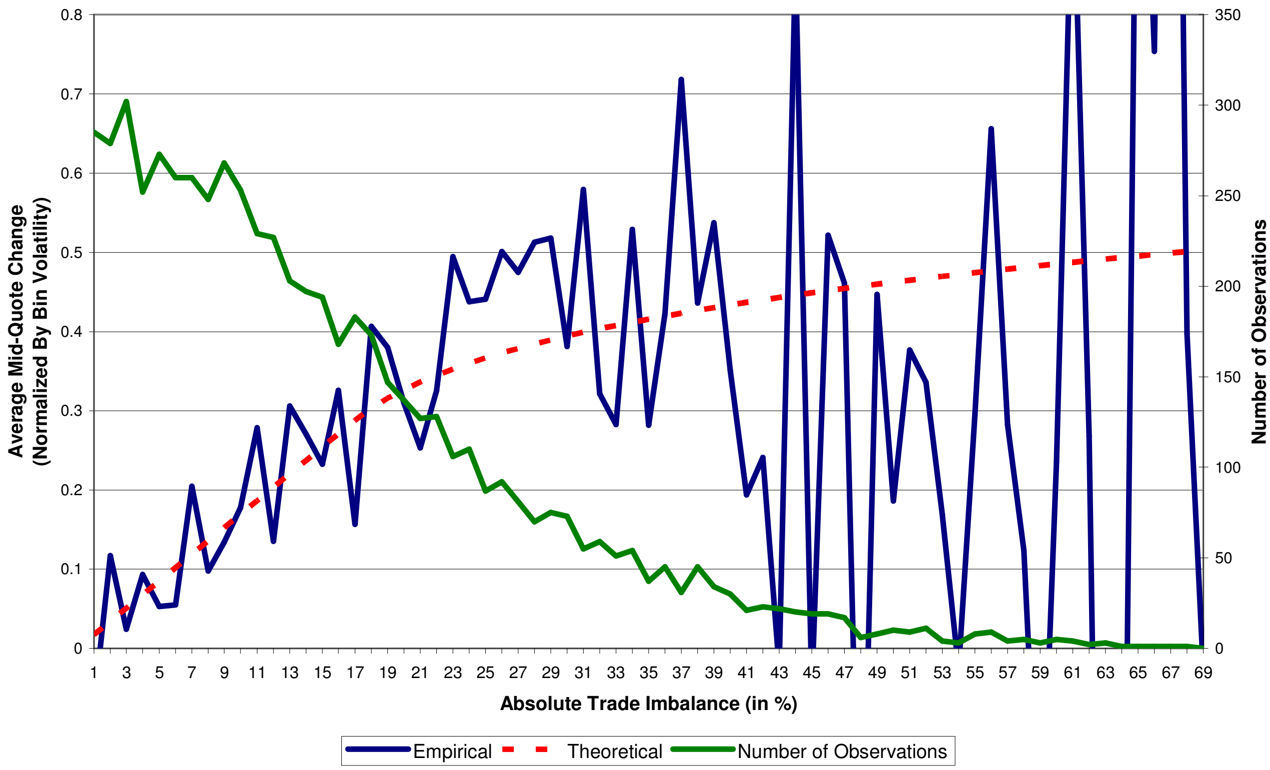


This figure shows the empirical permanent price impact function in bin 1 (9:30 a.m. - 10:00 a.m.) for the most liquid, U.S. Listed stocks (solid line). The empirical permanent price impact function is obtained by segmenting the observations in trade imbalance groups and then taking averages in each group. The empirical permanent price impact is linear until some point when it becomes concave. This behavior is the same for all time intervals, liquidity groups, and markets and can be observed for both permanent and temporary price impacts. Consequently, all theoretical price impact functions in ITG ACE are characterized by three parameters: the slope *s*, the value *x* that represents the order size at which concavity starts, and the concavity parameter *alpha*. The dashed line shows the fitted theoretical permanent price impact function.

**FIGURE 11**

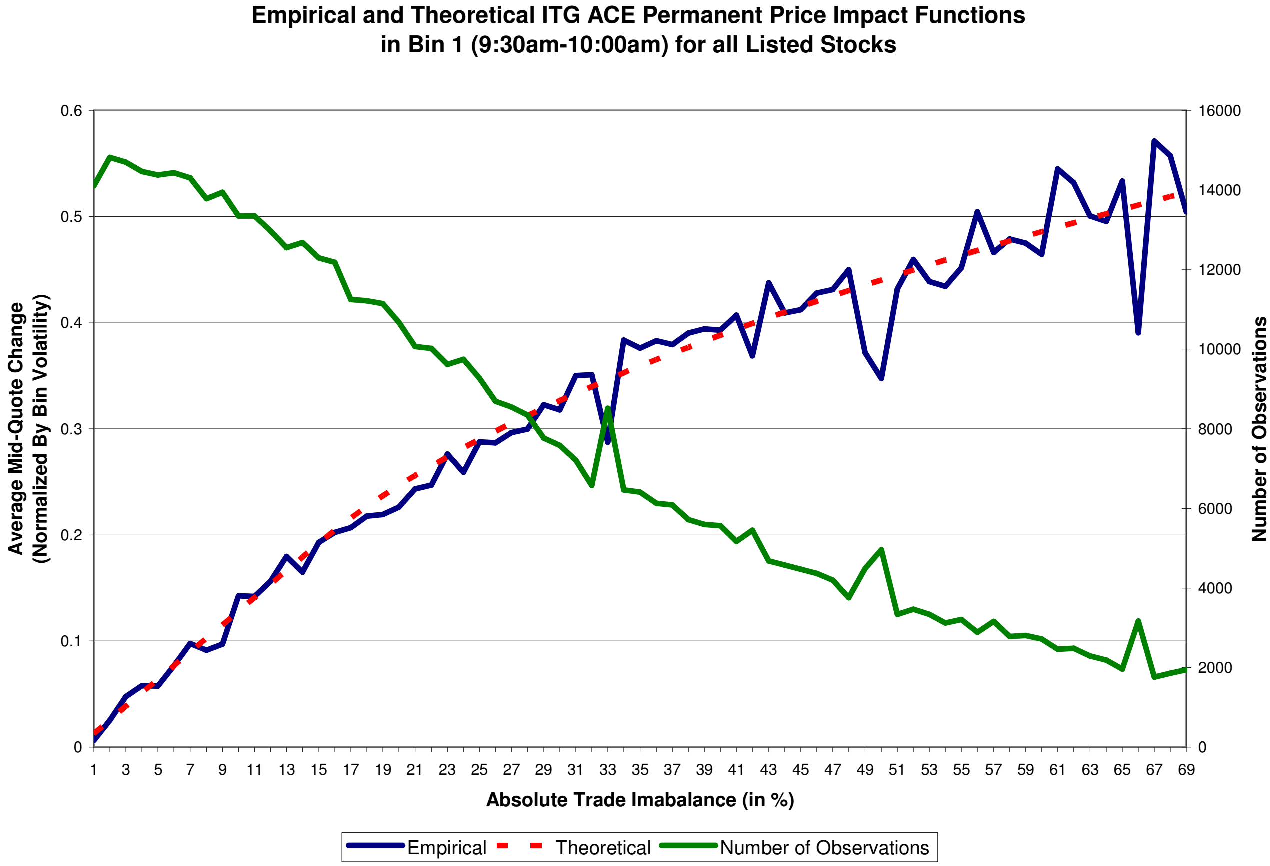


This figure shows the empirical permanent price impact function in bin 1 (9:30 a.m. - 10:00 a.m.) for the all U.S. Listed stocks (solid line). The empirical permanent price impact function is obtained by segmenting the observations in trade imbalance groups and then taking averages in each group. The empirical permanent price impact is linear until some point when it becomes concave. This behavior is the same for all time intervals, liquidity groups, and markets and can be observed for both permanent and temporary price impacts. Consequently, all theoretical price impact functions in ITG ACE are characterized by three parameters: the slope *s*, the value *x* that represents the order size at which concavity starts, and the concavity parameter *alpha*. The dashed line shows the fitted theoretical permanent price impact function. Compared to Figure 10, the empirical permanent price impact function is much smoother due to the aggregation over all U.S. Listed stocks.

**FIGURE 12**

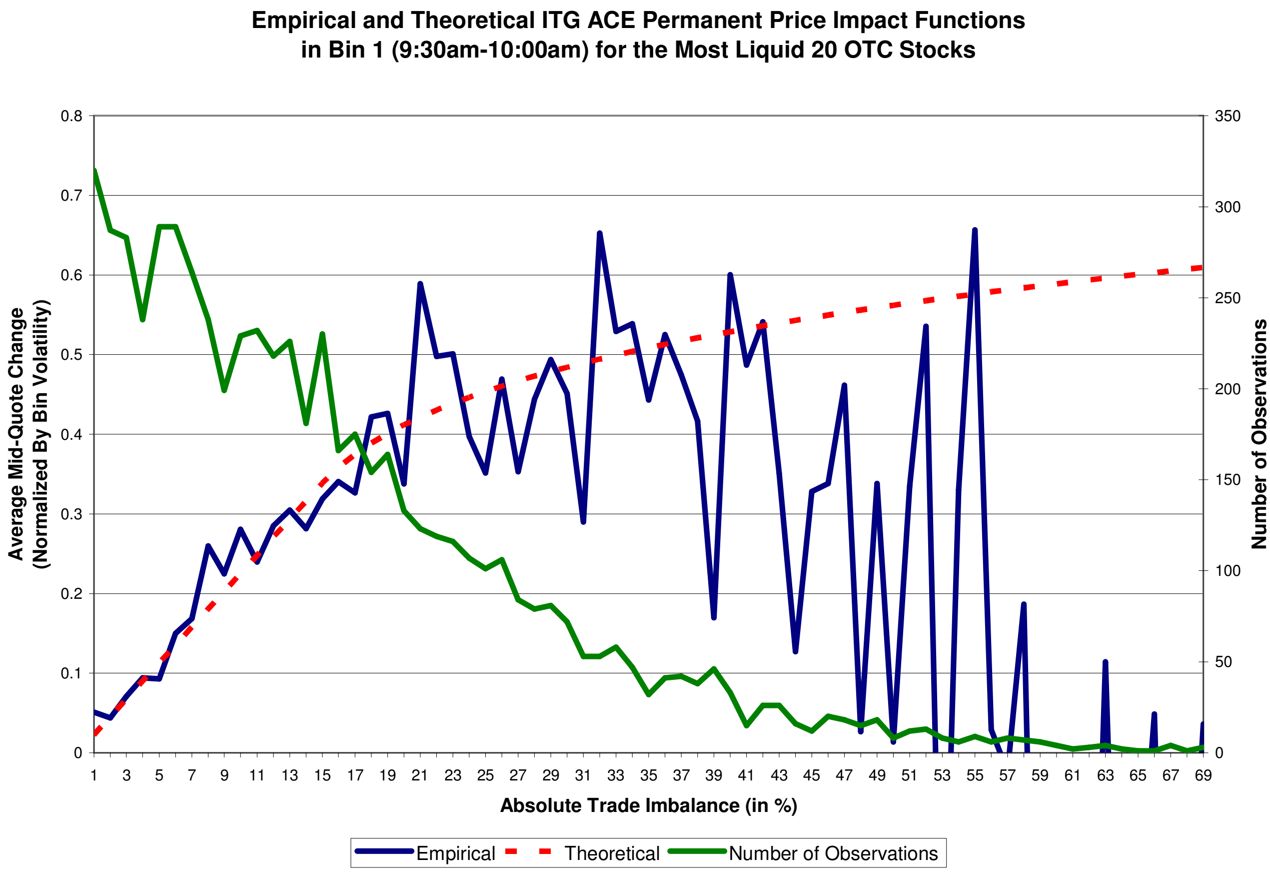


This figure shows the empirical permanent price impact function in bin 1 (9:30 a.m. - 10:00 a.m.) for the most liquid, U.S. OTC stocks (solid line). The empirical permanent price impact function is obtained by segmenting the observations in trade imbalance groups and then taking averages in each group. The empirical permanent price impact is linear until some point when it becomes concave. This behavior is the same for all time intervals, liquidity groups, and markets and can be observed for both permanent and temporary price impacts. Consequently, all theoretical price impact functions in ITG ACE are characterized by three parameters: the slope *s*, the value *x* that represents the order size at which concavity starts, and the concavity parameter *alpha*. The dashed line shows the fitted theoretical permanent price impact function.

**FIGURE 13**

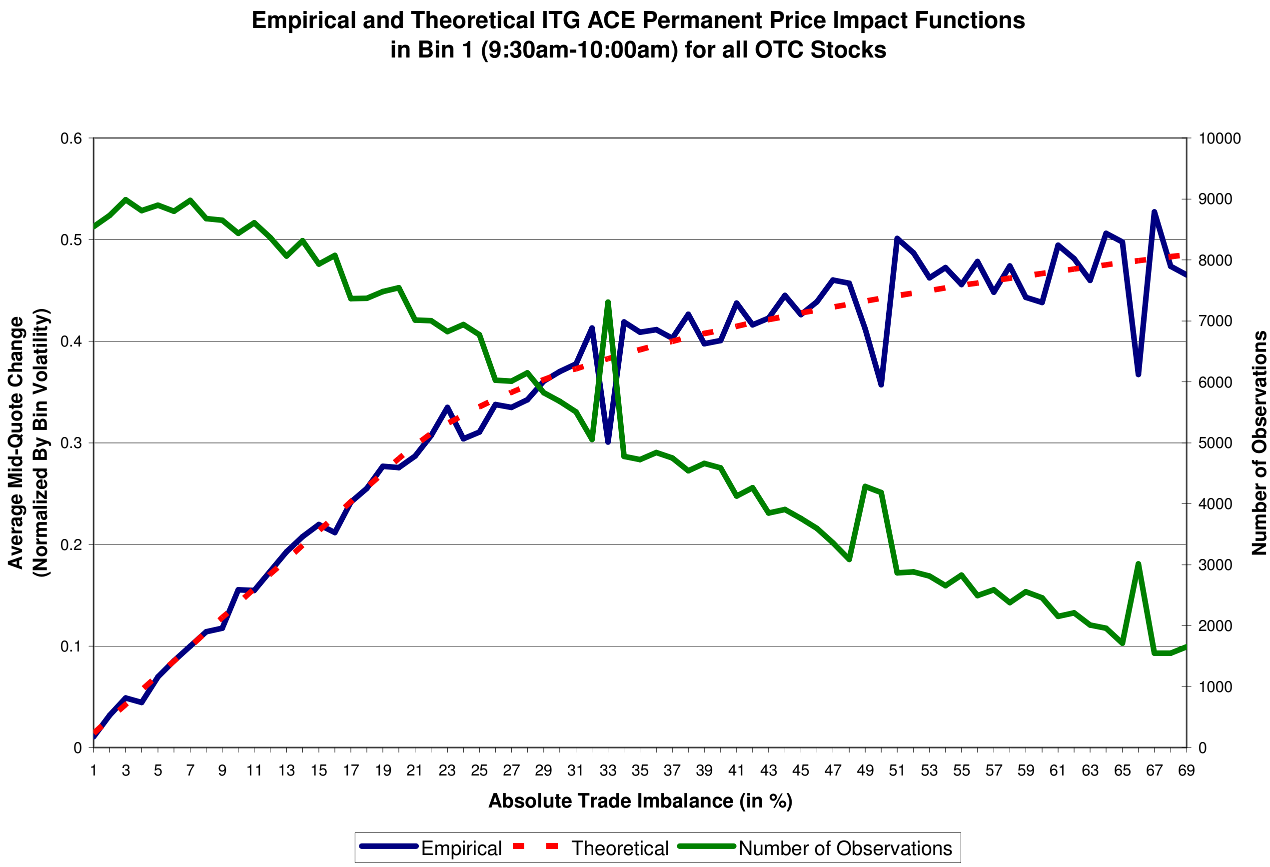


This figure shows the empirical permanent price impact function in bin 1 (9:30 a.m. - 10:00 a.m.) for the all U.S. OTC stocks (solid line). The empirical permanent price impact function is obtained by segmenting the observations in trade imbalance groups and then taking averages in each group. The empirical permanent price impact is linear until some point when it becomes concave. This behavior is the same for all time intervals, liquidity groups, and markets and can be observed for both permanent and temporary price impacts. Consequently, all theoretical price impact functions in ITG ACE are characterized by three parameters: the slope *s*, the value *x* that represents the order size at which concavity starts, and the concavity parameter *alpha*. The dashed line shows the fitted theoretical permanent price impact function. Compared to Figure 10, the empirical permanent price impact function is much smoother due to the aggregation over all U.S. OTC stocks.

**FIGURE 14**

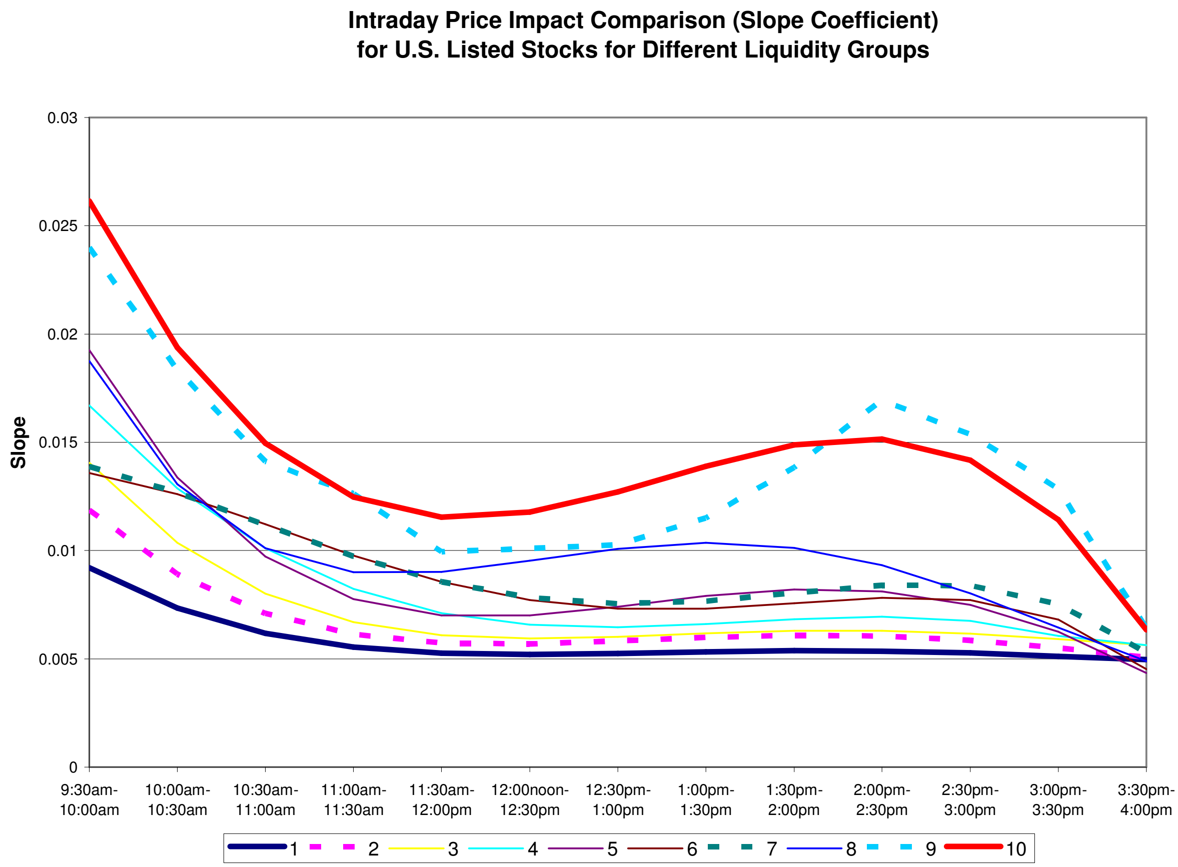


This figure illustrates the intraday pattern of the slopes of the permanent price impact functions for U.S. Listed stocks. The stocks are segmented into ten different liquidity groups. Stocks in all liquidity groups show the same intraday pattern. The price impact is the largest in the morning and is relatively low around noon and at the close.

**FIGURE 15**

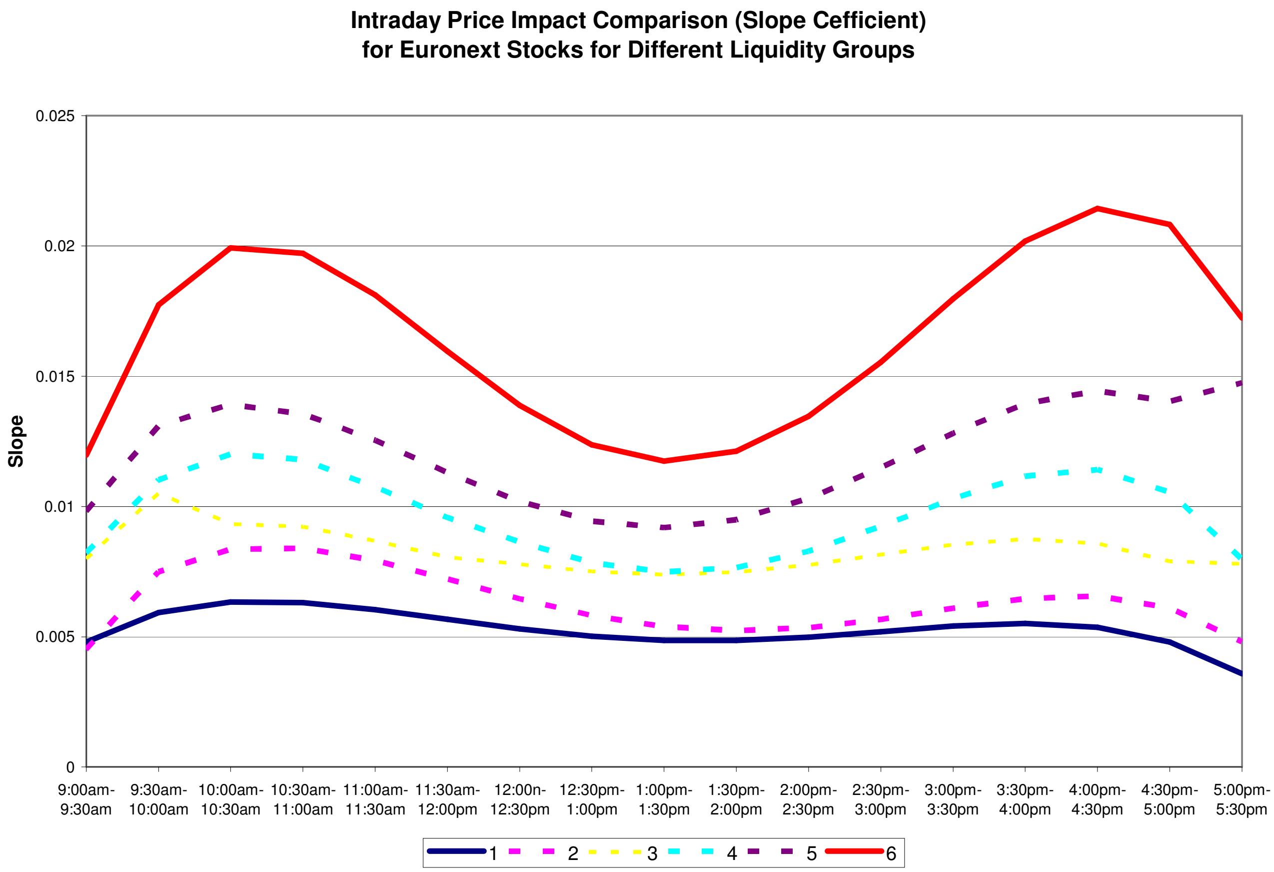


This figure illustrates the intraday pattern of the slopes of the permanent price impact functions for Euronext stocks. Euronext is the combined market of France, Belgium, Netherlands and Portugal. The stocks are segmented into six different liquidity groups, with Liquidity Group 1 being the least and Liquidity Group 6 being the most liquid stocks. Stocks in all liquidity groups show the same intraday pattern. The price impact is small in the morning, around noon, and at the close.

**FIGURE 16**

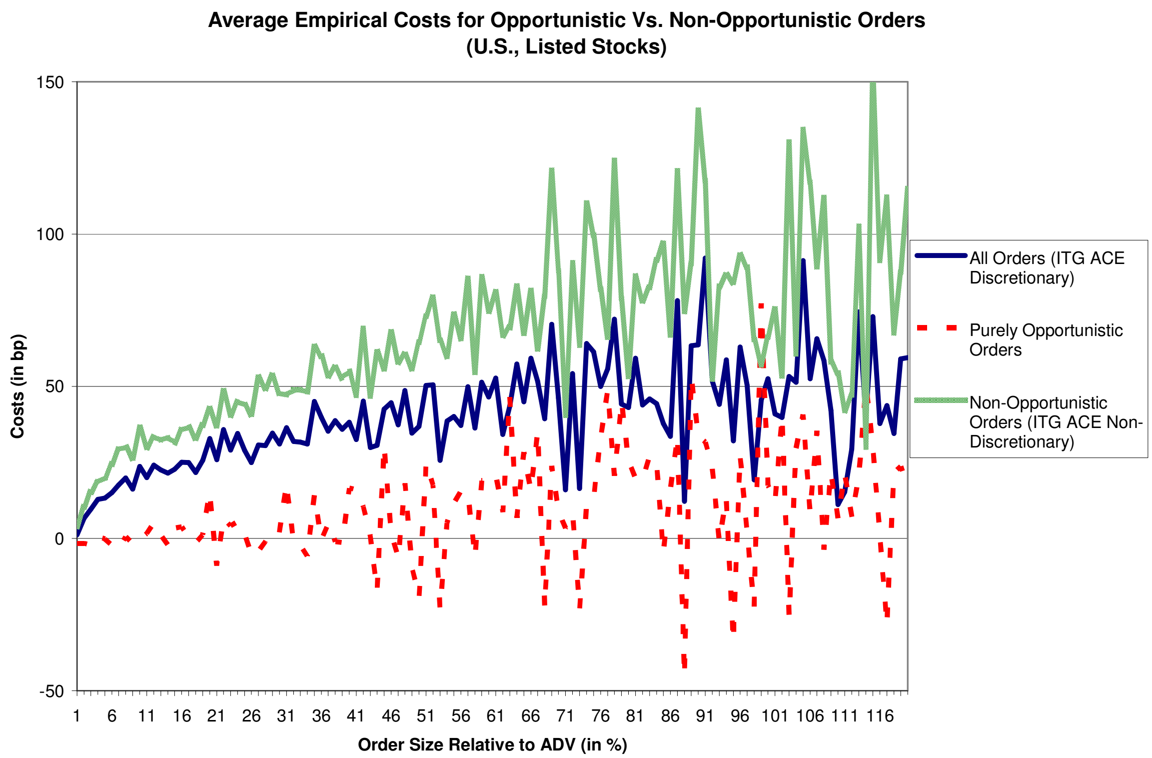


**FIGURE 17**

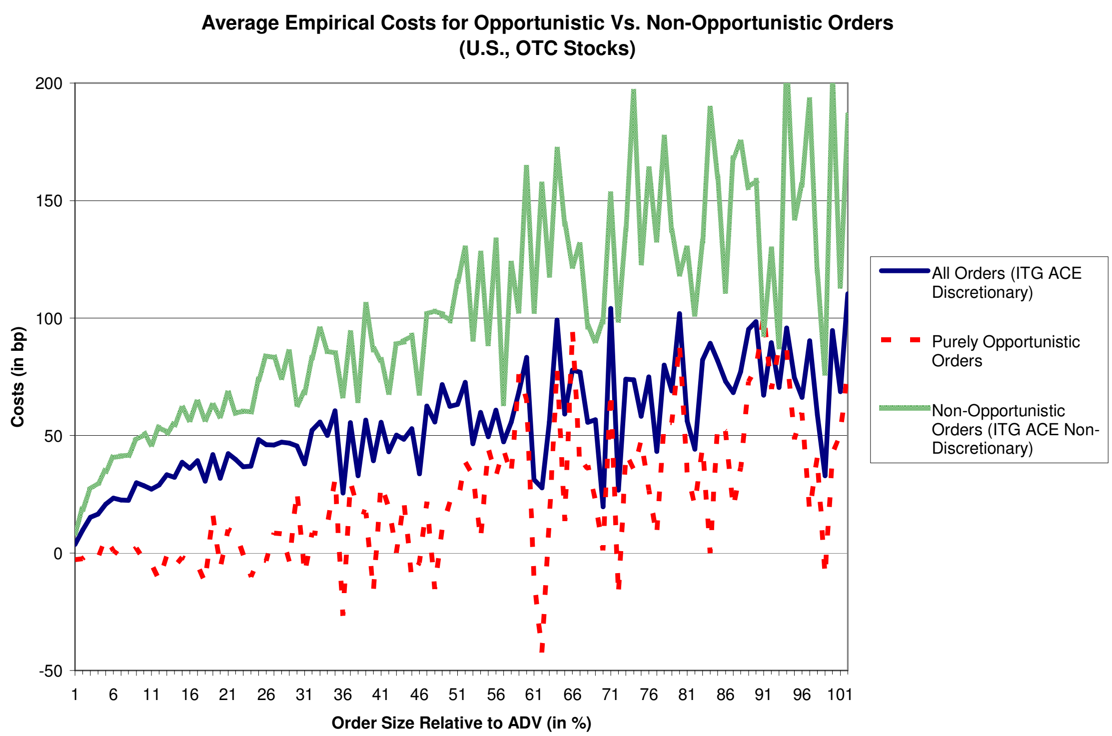


Figure 16 and Figure 17 plot the average empirical transaction costs that are associated with ITG ACE/2 Discretionary and ITG ACE/2 Non-Discretionary along with the average empirical transaction costs for purely opportunistic orders for Listed and OTC stocks in the U.S., respectively. The underlying execution data is the ITG Peer Group Database during the period from January 2005 to December 2006.

 

**FIGURE 18**

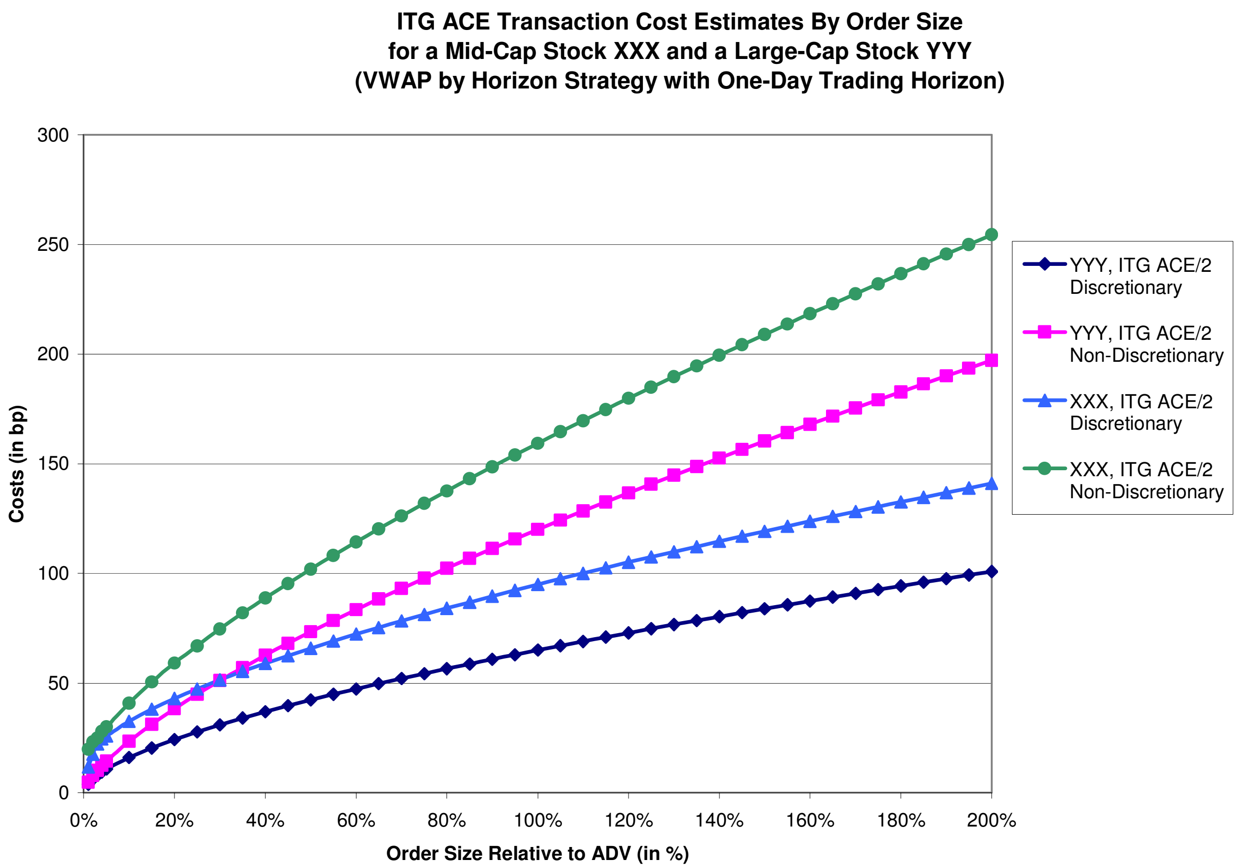


This figure shows the difference in transaction cost estimates for stocks XXX and YYY using ITG ACE/2 Non-Discretionary and ITG ACE/2 Discretionary. Stock XXX is a randomly chosen mid-cap stock and is a relatively illiquid stock; its market capitalization is about $390 million and the median daily share volume is about 50,000 shares. Stock YYY is a randomly chosen large-cap stock and is a relatively liquid stock; its market capitalization is over $70 billion and the median daily share volume is about 3.5 million shares. The ITG ACE cost estimates have been computed using a VWAP by Horizon Strategy with one-day trading horizon. The chart shows that transaction costs for orders that need to be completed are higher than those that reflect a market average amount of opportunistic trading.

**FIGURE 19**

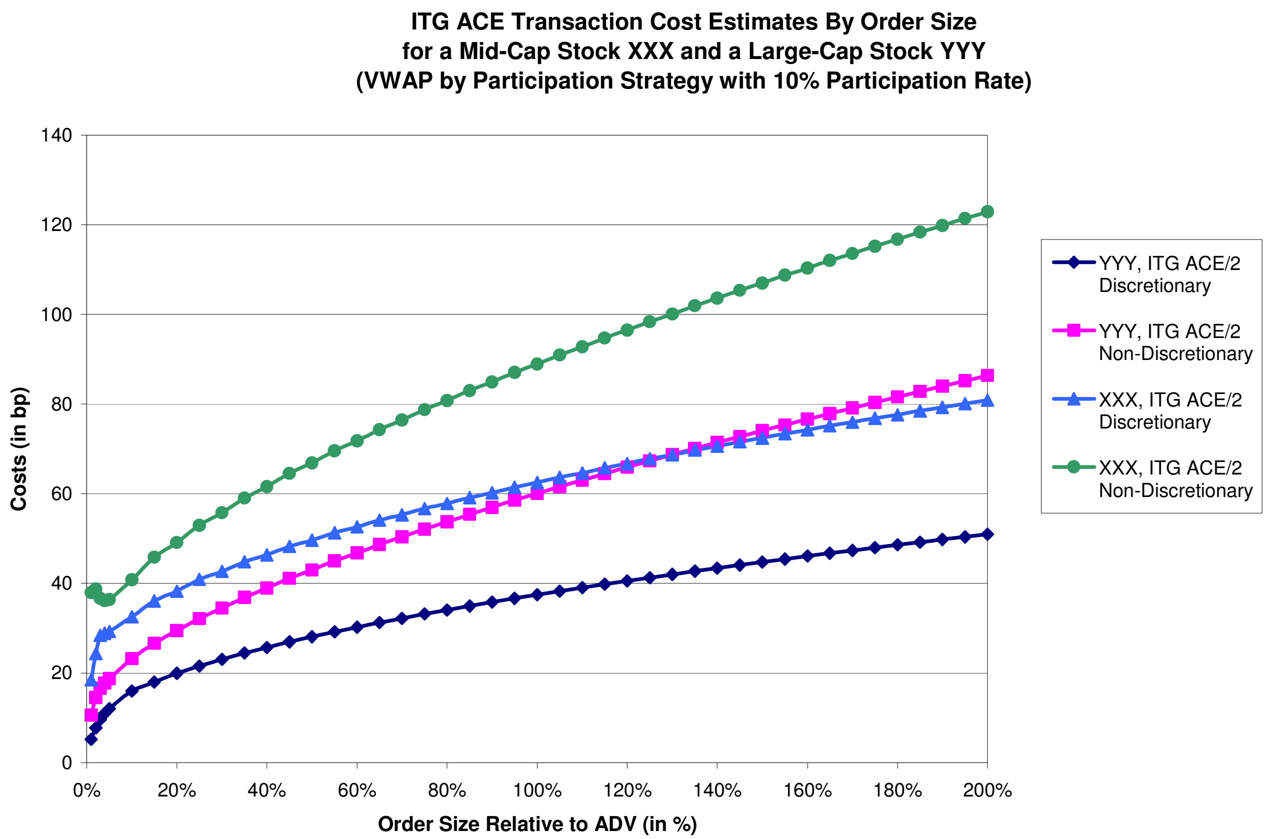


This figure shows the difference in transaction cost estimates for stocks XXX and YYY using ITG ACE/2 Non-Discretionary and ITG ACE/2 Discretionary. Stock XXX is a randomly chosen mid-cap stock and is a relatively illiquid stock; its market capitalization is about $390 million and the median daily share volume is about 50,000 shares. Stock YYY is a randomly chosen large-cap stock and is a relatively liquid stock; its market capitalization is over $70 billion and the median daily share volume is about 3.5 million shares. The ITG ACE cost estimates have been computed using a VWAP by Participation Strategy with 10% participation rate. The chart shows that transaction costs for orders that need to be completed are higher than those that reflect a market average amount of opportunistic trading.

**FIGURE 20**

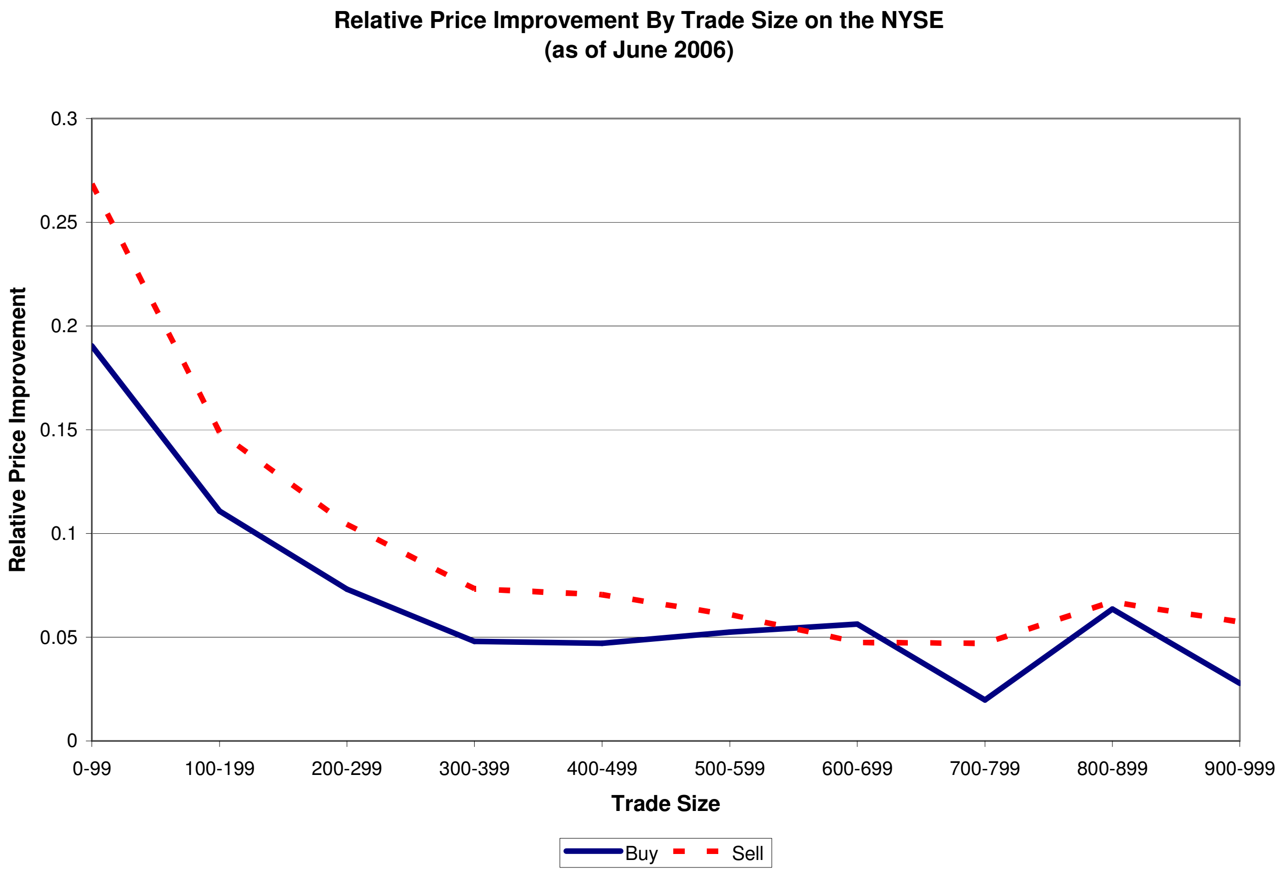


This figure reports the average empirical relative price improvement for stocks traded on the NYSE depending on trade size and trade side. Relative price improvement is defined in Equation (4) in Section 7 of this document. It lies between 0 and 1, with 0 indicating no price improvement and 1 indicating an execution at the other side of the spread. The graph is based on ITG proprietary execution data for June, 2006. The highest average relative price improvement can be observed for orders in size of less than 100 shares. The larger the order size, the less average relative price improvement can be observed. On average, sell trades get more price improvement than buy trades.

**FIGURE 21**

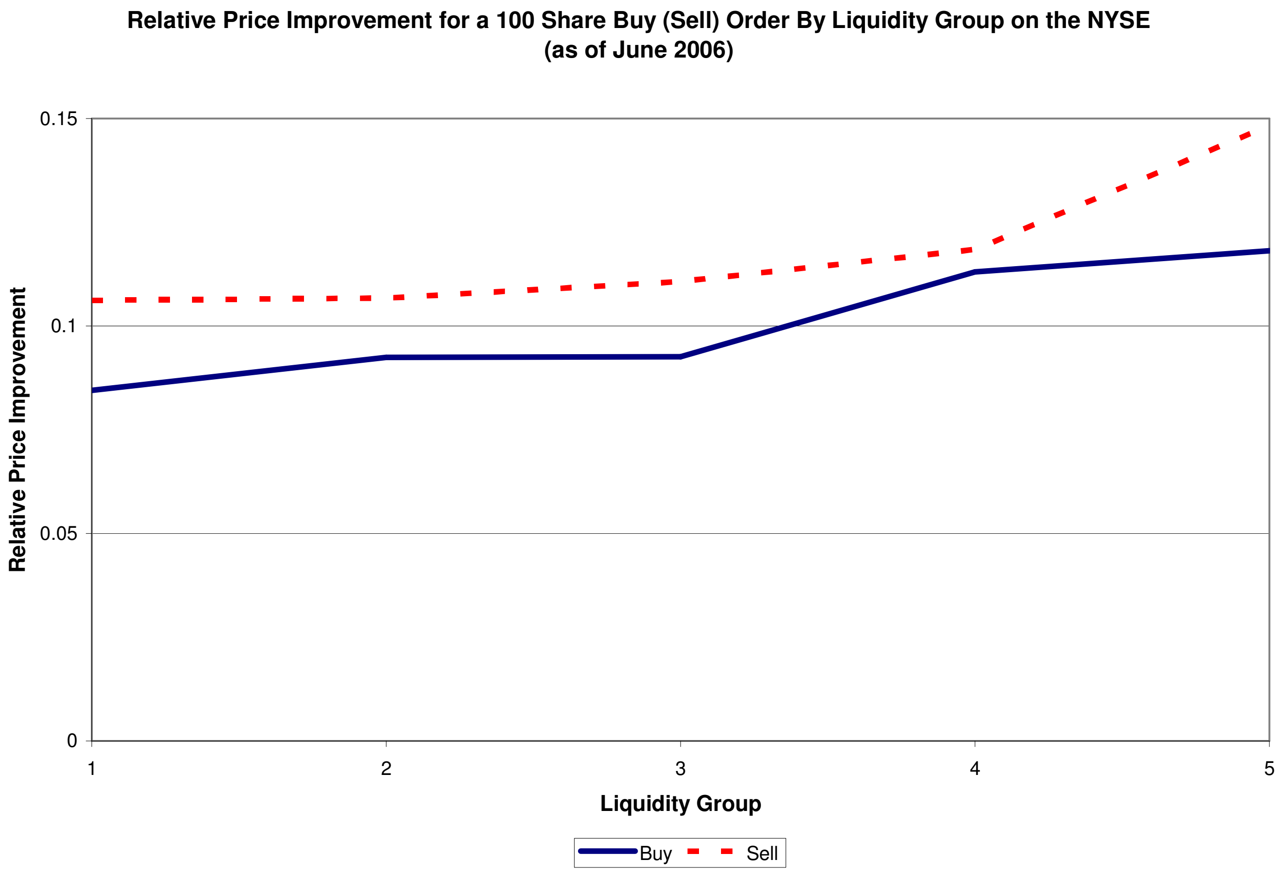


This figure reports average empirical relative price improvement for stocks traded on the NYSE that belong to different liquidity groups, with Liquidity Group 1 comprising the most and Liquidity Group 6 comprising the least liquid stocks. Relative price improvement is defined in Equation (4) in Section 7 of this document. It lies between 0 and 1, with 0 indicating no price improvement and 1 indicating an execution at the other side of the spread. The graph is based on ITG proprietary execution data for June, 2006. There is an almost linear relation between average price improvement and liquidity. Relative price improvement is the lowest for the most liquid stocks and the highest for illiquid stocks. On average, sell trades get more price improvement than buy trades.

**FIGURE 22**

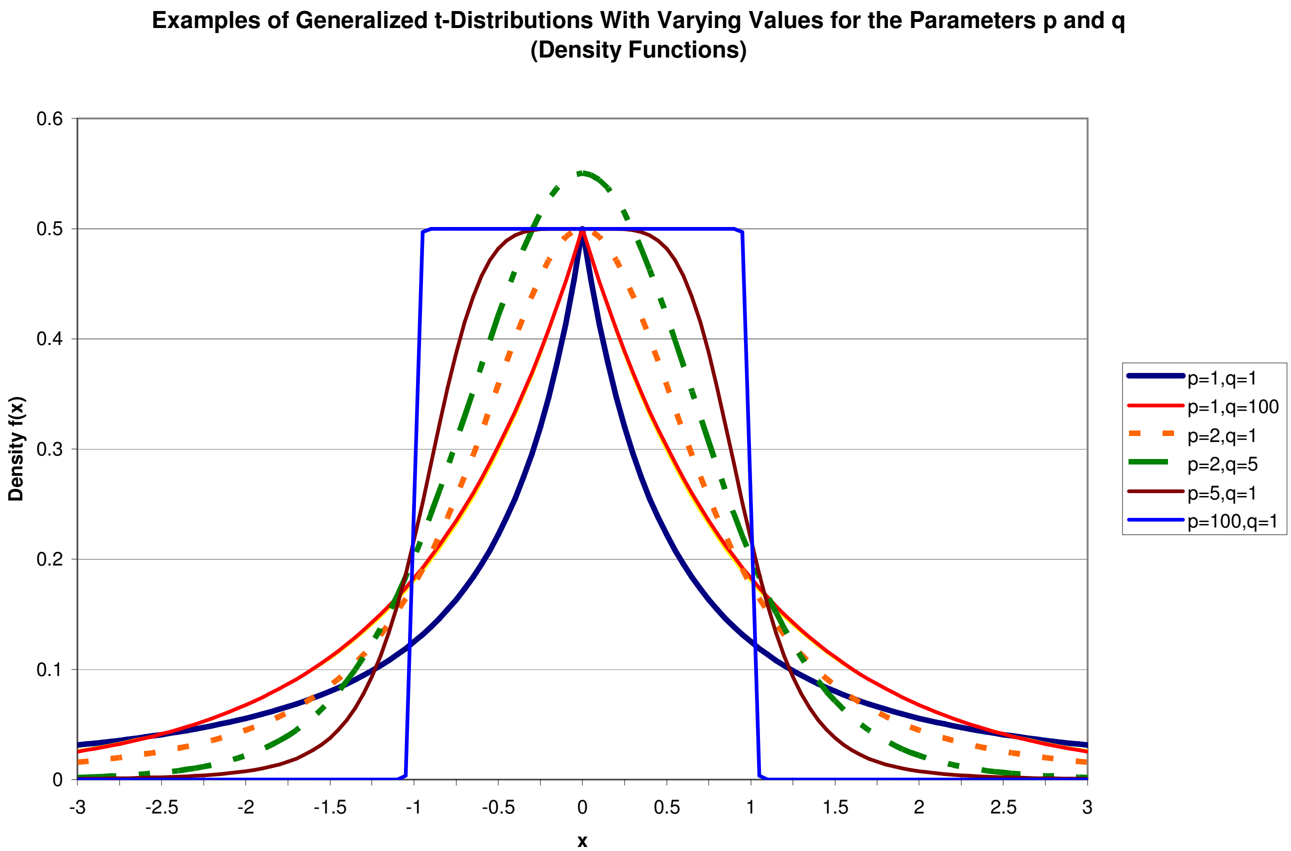


This figure illustrates different generalized t-distributions given the choice of the parameters $p$ and $q$. It is well-known that one can obtain the regular Student's t-distribution by setting $p=2$. As a consequence, $p=2$ and $q \rightarrow \infty$ yield the normal distribution.

 

**FIGURE 23**

**Comparison of Different Calibrated Distributions for the Empirical Costs z-Score**

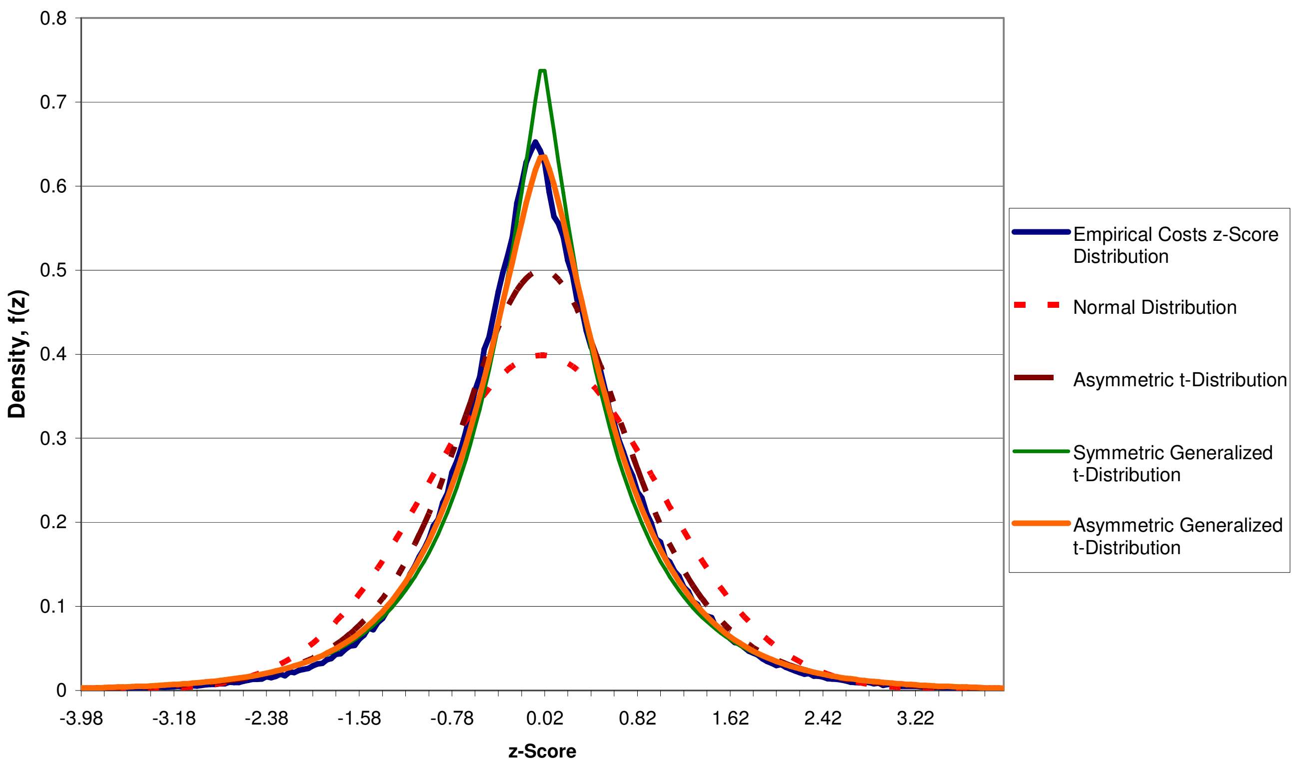


This figure compares the aggregated distribution of the z-scores of actual peer group database costs with four different calibrated distributions: the normal distribution, the asymmetric t-distribution, the symmetric generalized t-distribution and the asymmetric generalized t-distribution. Both the normal and the asymmetric t-distribution do not fit the empirical distribution well. The asymmetric generalize t-distribution captures all the observed properties. It is heavy-tailed, leptokurtic and asymmetric (the median is smaller than the mean).

# FIGURE 24

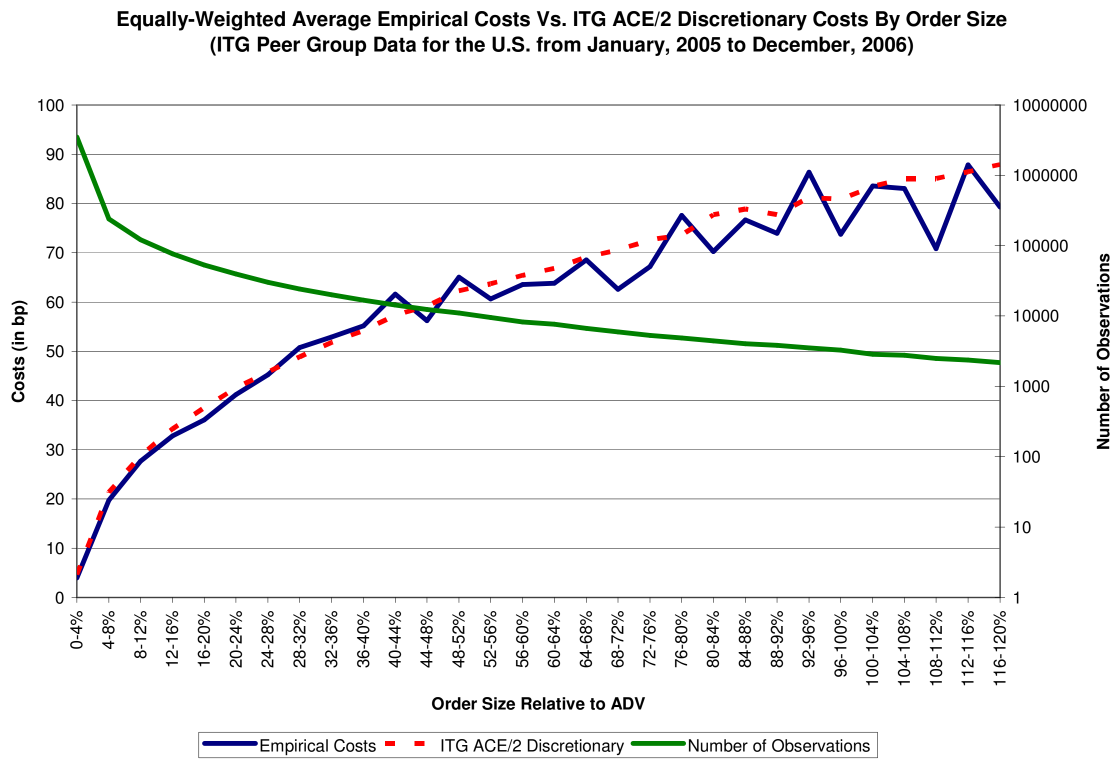


# FIGURE 25

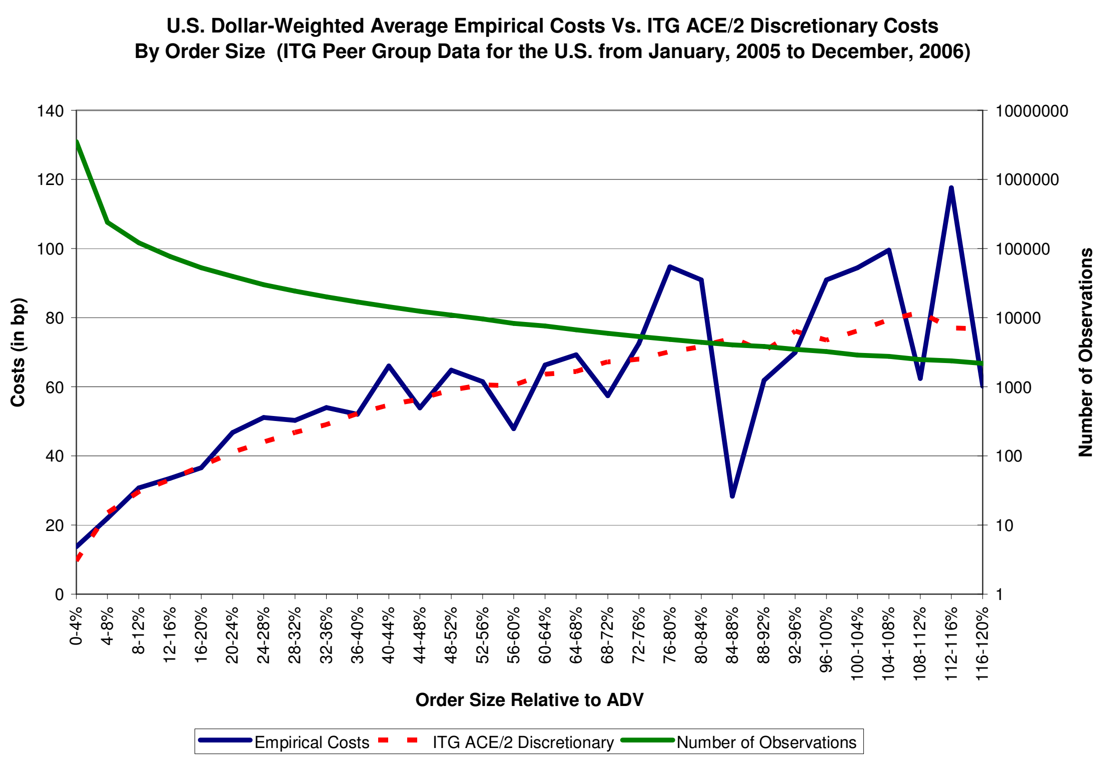


Figure 24 and Figure 25 show equally- and dollar-weighted average empirical costs and ITG ACE/2 Discretionary cost estimates for different order sizes for all U.S. trades in the ITG Peer Group Database from January 2005 to December 2006. The charts demonstrate a very good fit for the ITG ACE/2 Discretionary model. Similar fits can be observed for all other ITG ACE/2 countries and are available upon request.

## FIGURE 26

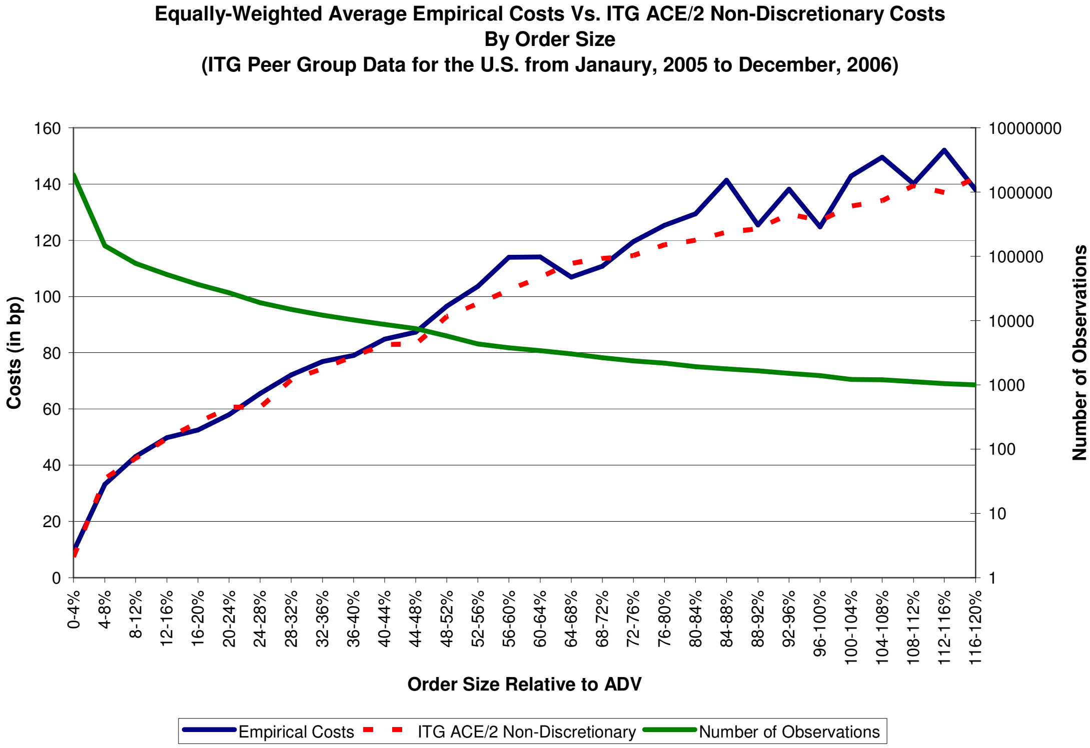


## FIGURE 27

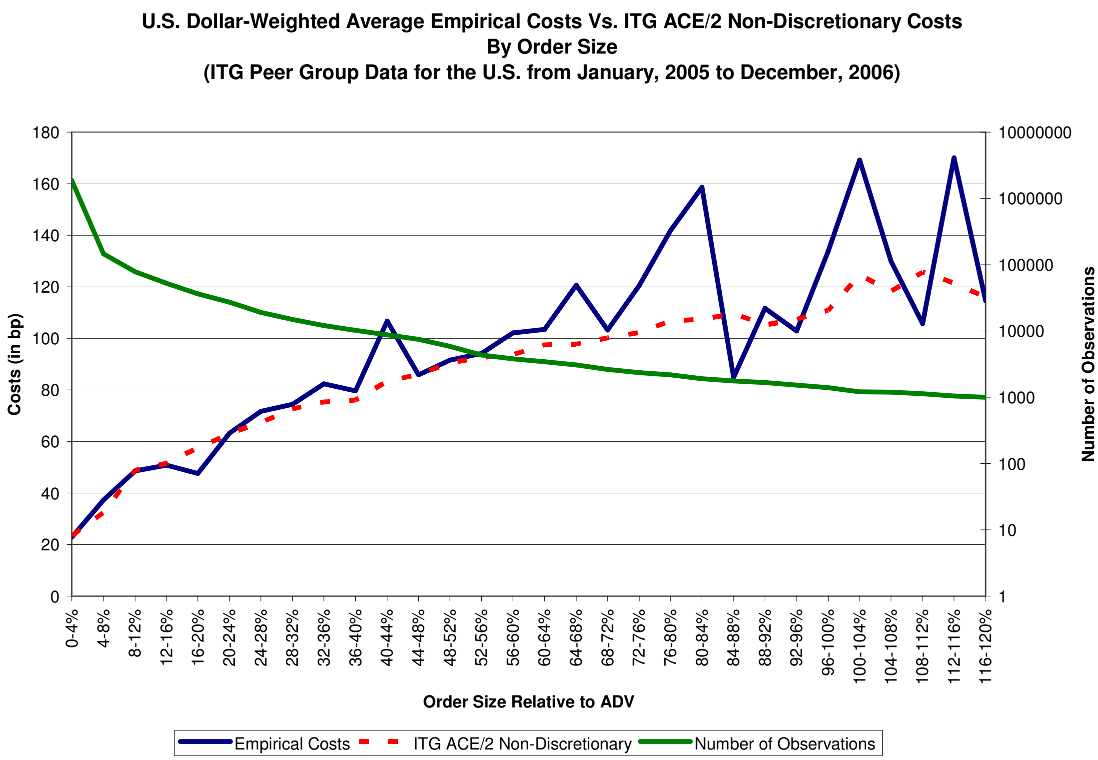


Figure 26 and Figure 27 show equally- and dollar-weighted average empirical costs and ITG ACE/2 Non-Discretionary cost estimates for different order sizes for all U.S. trades in the ITG Peer Group Database from January 2005 to December 2006. The charts demonstrate a very good fit for the ITG ACE/2 Non-Discretionary model. Similar fits can be observed for all other ITG ACE/2 countries and are available upon request.

## FIGURE 28

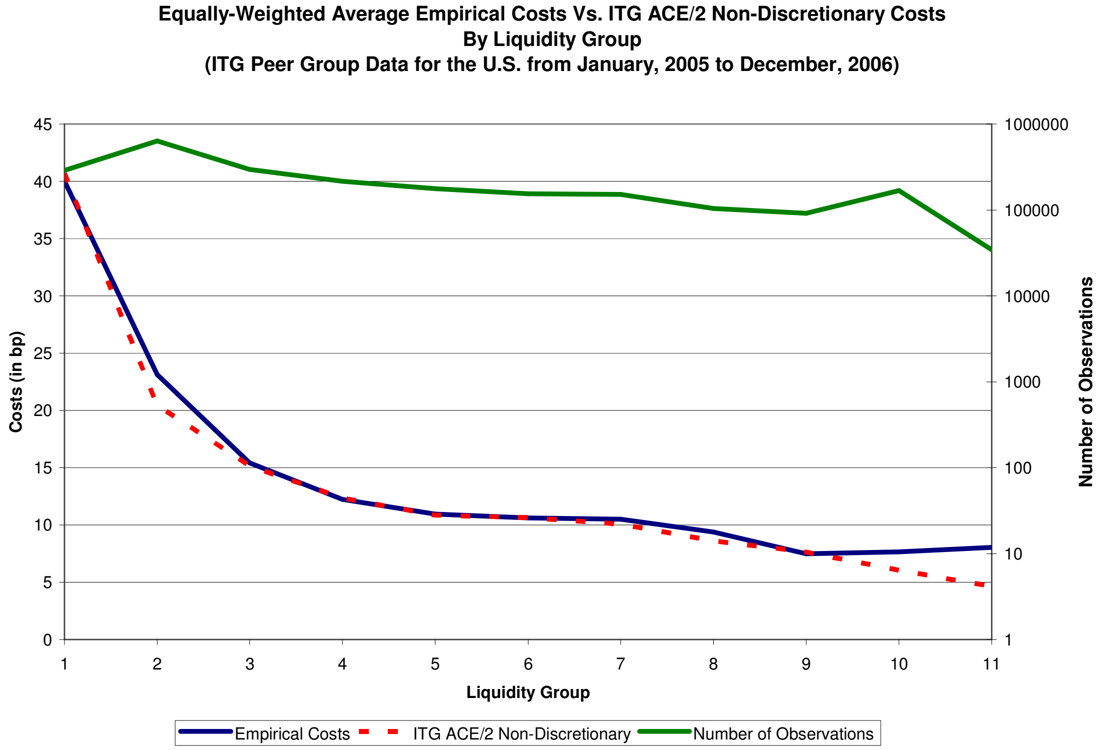


## FIGURE 29

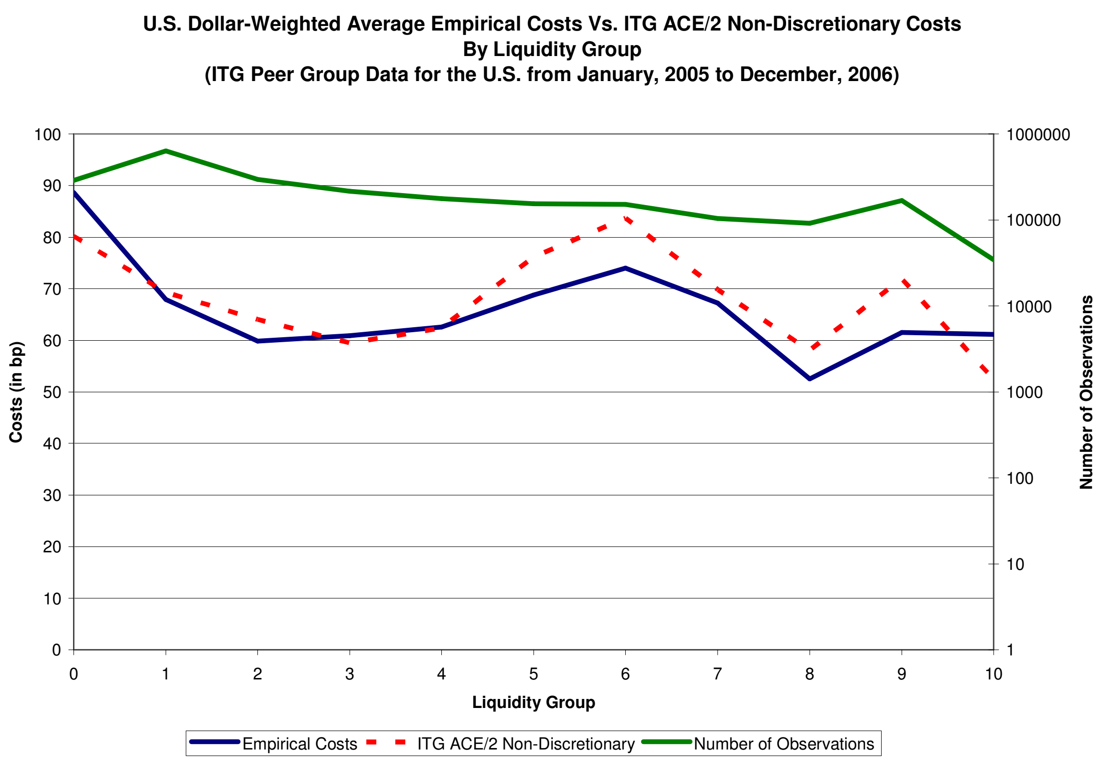


Figure 28 and Figure 29 show equally- and dollar-weighted average empirical costs and ITG ACE/2 Non-Discretionary cost estimates for different liquidity groups for all U.S. trades in the ITG Peer Group Database from January 2005 to December 2006. The charts demonstrate a very good fit for the ITG ACE/2 Non-Discretionary model. (Note: The higher the liquidity group number the more liquid the stocks are.)