Market Microstructure Invariants *

Albert S. Kyle
Robert H. Smith School of Business
University of Maryland
akyle@rhsmith.umd.edu

Anna A. Obizhaeva
Robert H. Smith School of Business
University of Maryland
obizhaeva@rhsmith.umd.edu

Preliminary Version - Please Do Not Circulate

October 31, 2009

Abstract

A simple theoretical model of market microstructure invariants is developed to generate hypotheses concerning how market depth, bid-ask spread, and order size vary across stocks with different level of trading activity. The model is tested using a dataset of portfolio transitions containing over 400,000 orders in individual stocks executed during the period 2001-2005. In a framework like Kyle (1985), our proposed model of "trading game invariance" assumes that the expected number and size of trades per "trading game" are invariant across stocks and across time, in contrast to alternative models which assume that the length of the "trading day" is invariant (e.g., equal to precisely one calendar day for all stocks). Our assumptions find a strong support in portfolio transition data. The proposed model predicts that for every one percent increase in the product of dollar trading volume with return volatility, i.e. the measure of trading activity, the market impact of trading one percent of average daily volume increases by one-third of one percent. Using implementation shortfall to estimate market impact in a non-linear regression, the parameter predicted to be one-third is estimated to be 0.33 with standard error of 0.024 (t-statistics of 13.37). Our model makes similar predictions about bid-ask spreads that find statistical support from regressions based on portfolio transition data as well. The proposed model implies simple formulas for market impact and bid-ask spread as functions of observable dollar trading volume and volatility.

*We are grateful to Georgios Skoulakis, Jim Gatheral, Mark Loewenstein, and Vish Viswanathan for helpful comments. Obizhaeva is also grateful to the Paul Woolley Center at London School of Economics for its hospitality as well as Ross McLellan, Simon Myrgren, Sebastien Page, and especially to Mark Kritzman for their help. Kyle is primarily responsible for the theoretical model in this paper. Obizhaeva is responsible for the empirical implementation. Both authors contributed equally to this paper.

1 Introduction

What is the appropriate benchmark for assessment of execution quality? How to compare performance of brokers executing trades in different stocks? How much money will be lost during changes in asset allocation or replacement of fund management? What percentage of "alpha" will be lost due to transaction costs during implementation of a seemingly profitable strategy? How much money can be allocated to this strategy before increasingly high costs make it non-economical? Should the trading strategies be similar across different securities? This paper develops a cross-sectional model of transaction costs that finds a strong support in empirical data and thus can provide answers to these questions.

When portfolio managers trade stocks, they can be modeled as playing trading games. Since portfolio managers trade many different stocks, we can think of them as playing many different trading games simultaneously, a different game for each stock. A trading game in which an informed trader, liquidity traders (or noise traders), and market makers trade one common stock was first described in Treynor (1971) and later formalized in Kyle (1985). The specifics of trading games seem to vary significantly when these games are being compared across stocks, for example, across actively and inactively traded securities.

The purpose of this paper is to investigate, both theoretically and empirically, what features of these trading games remain invariant as games themselves vary across stocks with different levels of trading activity. Different assumptions about market microstructure invariants lead to different predictions about how market depth and bid-ask spreads vary across stocks. These assumptions and predictions are tested using a proprietary database of portfolio transitions provided by a leading vendor of portfolio transitions services. In a portfolio transition, an incumbent portfolio manager is replaced by a newly hired one. The transition manager replaces the incumbent's legacy portfolio with a new portfolio by selling a portfolio held by the incumbent manager and buying a portfolio chosen by the new manager. A skilled transition manager tries to minimize the transactions costs, both market impact and bid-ask spread, associated with the transactions necessary for effecting the portfolio transition. Thus, a transition manager can be modeled as a liquidity trader who participates in trading games in many different stocks simultaneously over time.

Our proposed theory of "trading game invariance" is based on the idea that key features of market microstructure remain invariant when the trading games are compared across stocks and across time. These "deep" parameters include variables such as number of noise traders, size of risks taken by noise traders, size of risk taken by informed traders, profitability of private information, and the cost of acquiring this information per "trading game". Our theoretical model pays special attention to the frequency and size of liquidity trades, which we call "bets." The number of bets per trading game and the amount of risk transferred per bet is assumed to remain invariant across stocks. This assumption is consistent with the intuition that stocks differ only in the speed with which their trading games are played but the structure of trading games themselves is the same for each stock. Likewise, fast chess players play games with the same rules as normal chess tournaments but each side is given less time to make their moves.

We can define a measure of daily "trading activity," which we denote as W, as the product of dollar trading volume per calendar day and daily standard deviation of the stock's returns. According to this measure, active stocks are stocks with high volatility and high dollar trading volume per calendar day, while inactive stocks are stocks with low volatility and low dollar trading volume per calendar day. Our assumption implies that the trading games for active stocks and inactive stocks are the same, i.e., their deep parameters remain constant across stocks, per "trading game". The trading games for active stocks are, however, played at a faster pace than those for inactive stocks. This leads to the intuition that the length of a "trading day" differs from the length of a calendar day, with the trading day for active stocks perhaps corresponding to a few minutes while the trading day for inactive stocks perhaps corresponding to a few months. The length of the trading day is related to the market efficiency. The shorter is the trading day, the more efficient is the market.

The invariance of trading games (i.e. of bet sizes and bet frequencies) across stocks, despite the difference in their measures of trading activity, is a reasonable assumption. For example, compare a liquidity trade in an active stock with a liquidity trade in an inactive stock of equal returns volatility. Our measure of bet size is a product of the stock price, the number of shares traded, the daily percentage standard deviation of the stock's return, and the square root of the length of the trading day. This measure captures the amount of risk transferred by the trade, taking account of the trading horizon, which is assumed to be proportional to the length of the trading day. The dollar trade size is larger for the active stock than for the inactive one. The trading day for the active stock is, however, shorter. A larger position held for a shorter period of time can have the same risk as a smaller position held for a longer period of time. Both transactions can represent the same amount of risk transferred. In the context of our model, the expected amount of risk transfer per liquidity trade can be invariant across the two markets. Also, compare the number of bets in an active stock with the number of liquidity trades in an inactive stock. The number of bets per trading game is equal to the number of bets per calendar day multiplied by the length of the trading day. The active stock has more bets taking place over one calendar day than the inactive stock. The trading day for the active stock is, however, shorter than for the inactive stock. In the context of our model, the number of bets per trading game can be invariant across the two markets.

If the only difference between market microstructure of active and inactive stocks is the speed with which the time passes, this implies a particular way in how size of liquidity trades and their frequency should vary with trading activity. Our model assumes that if trading activity W increases by one percent, then one-third of this increase comes from increased bet size per calendar day and two-thirds of this increase results from increased bet frequency per calendar day. Consequently, the size of liquidity trades as a fraction of daily trading activity has to decrease by two-thirds of one percent per a one percent increase in trading activity (holding returns volatility constant). Our tests based on the portfolio transition data validate these assumptions.

What is the intuition for the "one-third" and "two-third" fractions appearing in our assumptions? These fractions result from the fact that bet size and bet frequency do not change with the time horizon in the same manner: Being a measure of risk transfer, bet size is proportional to the square root of the time horizon whereas bet frequency is proportional to the time horizon itself. Consider the following example. The trading activity increases by one percent, bet size and bet frequency per calendar increase as assumed by our model. Suppose also that the length of the trading day simultaneously decreases by two-thirds of one percent, reflecting a faster time pace. Then the number of bets per trading day is constant.

Although the bets are larger by one-third of a percent, they are held for a length of time two-thirds of one percent shorter. This makes the riskiness of the bets constant as well, since the riskiness of a bet is proportional to its size and to the square root of the time for which it is held. Another way to develop the same intuition is to ask what happens when the length of the trading day is shortened. For the number of bets per trading game to remain constant, the rate at which bets are made must increase proportionately. For these bets to be equally risky, the size of the bets needs to increase only half as fast as the trading day is shortened because riskiness is proportional to the standard deviation of price changes, not to variance, i.e. to the square root of the time for which it is held.

Our theoretical model leads to predictions concerning how the magnitude of market impact and bid-ask spread varies as a function of trading activity across different stocks. Our theoretical model predicts that a one percent increase in trading activity W leads to an increase of one-third of one percent in the market impact and to a decrease of one-third of one percent in the spread costs incurred in executing a liquidity trade equal to one percent of average daily volume, where transactions costs are measured in basis points per dollar traded (holding returns volatility constant). The prediction for price impact is derived from the formula for λ in Kyle (1985). We believe that other market microstructure models would lead to similar results. Bid-ask spread is assumed to be zero in Kyle (1985). Our prediction for bid-ask spread is derived based on an argument that the bid-ask spread should be inversely proportional to the price impact. This would be the case, for example, if market makers were not perfect competitors, as in Kyle (1983).

The predictions of our proposed model of "trading game invariance" are compared with the predictions of two alternative models based on different assumptions concerning market microstructure invariants. Both models are "naive" in the sense that they assume that a "trading day" is equivalent to one calendar day for all stocks. Both models are "extreme" in the sense that they assume that all variation in trading activity comes either entirely from variation in bet size or from variation in bet frequency.

The first alternative model assumes that as trading activity increases, the bet frequency per day remains invariant at some constant level, while the bet size varies proportionally with trading activity. We call this model - the model of "invariant bet frequency", since all variation in trading activity is explained entirely by variation in bet size. According to this model, trade size as a fraction of daily volume should remain invariant across stocks. Concerning market impact and spreads, this model predicts that as trading activity increases, the cost of executing a trade of one percent of average daily volume remains constant in basis points per dollar traded (holding returns volatility constant). We believe that this model is the "default model" that implicitly but incorrectly guides the intuition of many asset managers. For example, this model justifies trading of no more than say 5% of average daily volume for all stocks, regardless of their level of trading activity.

The second alternative model assumes that as trading activity increases, the bet size per calendar day remains invariant at some constant level, while the bet frequency per calendar day increases proportionately. We call this model - the model of "invariant bet size", since all variation in trading activity is explained entirely by variation in bet frequency.According to this model, trade size (adjusted for volatility) remains constant as trading activity changes, and therefore trade size as a proportion of daily volume falls at the same rate as trading activity rises. Concerning market impact and spreads, this model predicts that a one percent increase in trading activity leads to an increase of one-half of one percent in the price impact and to a decrease of one-half of one percent in the spread costs incurred in executing a liquidity trade equal to one percent of daily volume, measured in basis points per dollar traded (holding returns volatility constant).

These alternative models have simple intuition. In the model of Kyle (1985), market impact λ is proportional to the ratio of price volatility σ_V and the standard deviation of the inventories of noise traders σ_U, i.e., $\lambda = \sigma_V/\sigma_U$. To generate predictions about how market impact varies cross-sectionally as volume changes, it seems necessary to map σ_U into volume. One approach is to think of volume as proportional to the standard deviation σ_U. This is the approach taken by our first alternative model. Another approach is to think of volume as proportional to the variance σ_U^2. This is the approach taken by our second alternative model. In some sense, either of these two alternative choices is arbitrary. In our proposed model, the choice is made naturally by assuming that the trading game remains the same; it is only the speed with which the game is played that varies across stocks. Our model implies that volume per calendar day is proportional to $\sigma_{U,1}^{3/2}$ where $\sigma_{U,1}$ is the standard deviation of the inventories of noise traders over one calendar day.

Using portfolio transition data, the assumptions and predictions of three models are tested to examine which of them better describes the data. We exploit the data on portfolio transition orders to test the assumptions concerning the size of liquidity traders' bets, and we exploit the data on implementation shortfall of portfolio transitions to test the predictions concerning trading costs.

Our theoretical model as well as two alternative models assume that expected trade size should vary with daily trading activity W in a certain way. The predictions concerning average trade size $\bar{Q}$ as a fraction of volume V can be captured by the formula

$$\frac{\bar{Q}}{V} = \bar{q} \times \left[\frac{W}{(0.02)(40)(10^6)}\right]^{a_0}.$$

The model of trading game invariance predicts that a one percent increase in trading activity leads to a decrease of 2/3 of one percent in trade size as a fraction of daily volume. In the context of regressions, the above formula implies $a_0 = -2/3$. The model of invariant bet frequency implies $a_0 = 0$, and the model of invariant bet size implies $a_0 = -1/2$.

The assumptions of all three models are tested using portfolio transitions data. We make the identifying assumption that the size of portfolio transition trades is proportional to the size of liquidity trades in the theoretical model. Estimates of the above regression for trade size provide strong support for the model of trading game invariance. The coefficient estimate of -0.63 is remarkably close to the predicted value of $-2/3$. The assumptions of our model are also validated across different subsets of stocks.

Implementation shortfall is used to estimate market impact and bid-ask spread from the portfolio transition data. Perold (1988) defines implementation shortfall as the difference between a "paper trading" benchmark and actual trading results, marking to market unexecuted shares at post-trade prices. For our purposes, the paper trading benchmark for a given transition is defined to be the price which would have been obtained if all shares were executed at the market closing price the day before any trades implementing a given transition began to take place. This benchmark is compared against the actual prices at which the transition trades are later executed. The difference, measured in basis points per dollars worth of shares traded, measures implementation shortfall. Implementation shortfall has several components. It includes the effect of both market impact and bid-ask spread as well as random changes in the stock price between the benchmark date and the time when the trades are executed. We make the identifying assumption that the returns on the stock would otherwise have had a mean return of zero, which implies that the mean of the implementation shortfall is a measure of transactions costs.

There are two major problems usually associated with using implementation shortfall to estimate transactions costs. The database of portfolio transition data avoids both of these problems.

The first problem is low statistical power. The transaction costs are estimated using data on the realized implementation shortfall which is determined not only by transaction costs of a given trade but also by overall movements in stock prices. Since the latter contain significant amount of noise, as pointed out by Black (1986), even a properly specified regression to estimate transactions costs using implementation shortfall is going to have very low statistical power. The portfolio transition database addresses this problem in several ways. First, the data involves more than 400,000 individual orders executed over the period 2001-2005. The large number of degrees of freedom increases the statistical power of our estimates. The errors in regressions are potentially correlated due to the fact that many stocks are traded on the same days, and stock returns are correlated with one another. We pool observations at weekly levels for 17 industries. This pooling reduces degrees of freedom, but generates more accurate standard errors. Second, some of the orders are large enough to induce relatively significant market impact. This increases statistical power as well. As a result, our statistical tests are powerful enough to distinguish the proposed model from the two alternatives.

The second problem is potential selection bias. When transaction costs are estimated based on data sets containing prices and quantities of executed orders, these estimates are most likely biased downwards. The reason is that high cost orders are often canceled before execution and thus not observed in the data. Obizhaeva (2009) finds that the selection bias associated with unexecuted trades can be substantial. The data on portfolio transitions does not suffer from selection bias problem. The transition manager's job is to sell the entire legacy portfolio and replace it with the entire new portfolio. The lists of securities in both portfolios are specified in a transition mandate before trading. Assuming the transition manager executes each portfolio fully, the problem of selection bias due to unexecuted orders goes away.

Our theoretical model as well as two alternative models imply that market impact and bid-ask spread can be estimated from a non-linear regression in which the left-hand side is implementation shortfall of portfolio transition trades measured in basis points per dollar traded. In particular, our regression analysis is based on the formula for the expected trading costs, denoted $C(X)$, for an order of X shares

$$C(X) = \frac{1}{2}\bar{\lambda} \times \left(\frac{W}{(0.02)(40)(10^6)}\right)^{\alpha_0} \frac{\sigma_r}{0.02}\frac{X}{(0.01)V} + \frac{1}{2}\bar{k} \times \left(\frac{W}{(0.02)(40)(10^6)}\right)^{\alpha_1} \frac{\sigma_r}{0.02},$$

where trading activity W is the product of stock price P, trading volume V per calendar day, and daily returns volatility σ_r. There are two right-hand-side variables, one for market impact and one for bid-ask spread. The right-hand-side variable for market impact is transition order size X as a fraction of daily volume V. The coefficient for market impact is predicted to be proportional to a power of daily trading activity W. The coefficient associated with trade size can be thus written as $\frac{1}{2}\bar{\lambda}W^{\alpha_0}$. The right-hand-side variable for bid-ask spread is a constant term. The coefficient for spread is predicted to be proportional to a power of daily trading activity W. The coefficient associated with a constant term can be thus written as $\frac{1}{2}kW^{-\alpha_1}$. We define an arbitrary "benchmark stock" as a stock with a price of \$40 per share, trading volume of one million shares per day, and returns standard deviation of 2% per day. Market impact is scaled so that $\bar{\lambda}$ measures in basis points the market impact of trading one percent of the average daily volume in the benchmark stock. The market impact is multiplied by one-half because $\bar{\lambda}$ measures marginal price impact, but implementation shortfall captures average price impact, which is one-half marginal price impact. Spread is scaled so that $\bar{k}$ measures the bid-ask spread for the benchmark stock, measured in basis points. The bid-ask spread is multiplied by one-half because one-way trade incurs a spread cost of half the bid-ask spread $\bar{k}$.

In this cost formula, our proposed model of trading game invariance predicts that $\alpha_0 = 1/3$ and $\alpha_1 = -1/3$. Our two alternative models make different predictions. The model of invariant bet size predicts $\alpha_0 = \alpha_1 = 0$, while the model of invariant bet frequency predicts $\alpha = 1/2$ and $\alpha_1 = -1/2$.

The predictions of all three models are tested using portfolio transitions database. The model of trading game invariance predicts transactions costs from market impact and spread better than the other two alternatives. The empirical prediction that a one percent increase in trading activity increases the market impact (in units of daily standard deviation) by one-third of one percent is almost exactly the point estimate from non-linear regressions based on implementation shortfall. This provides strong support for our proposed model.

Our paper provides a simple formula for market impact and bid-ask spread as functions of observable dollar trading volume and volatility. Portfolio transition data allows us to estimate level of transaction costs for the benchmark stock. If the exponent parameters are fixed at levels implied by the model of trading game invariance, i.e., $\alpha_0 = 1/3$ and $\alpha_1 = -1/3$, then the estimated value of half market impact is $\lambda^*/2 = 2.89$, and the estimated formula for the half-spread is $k^*/2 = 7.91$; thus, a trade of one percent of average daily volume in the benchmark stock incurs a market impact cost of 2.89 basis points and a bid-ask spread cost of 7.91 basis points. Our proposed theoretical model predicts a certain cross-sectional variations in transaction costs. In particular, the above equation for $C(X)$ shows how to extrapolate the estimates for the benchmark stock to any other security.

Although our model of trading game invariance is based on the intuition that the trading day for active stocks is shorter than for inactive stocks, our data does not make it possible to identify the length of the trading game itself. To identify the length of the trading day, additional data would be needed such as, for example, data on the half-life of positions taken by traders. In fact, it is possible that, holding trading volume constant, the length of the trading day has been changing over time. For example, the increase in algorithmic trading may be associated with a shorter trading day.

Relations between many variables were found to follow power laws in nature, i.e., one variable is proportional to a power of another variable. Numerous power-law relations have been empirically established in physics, mathematics, linguistics, geophysics, sociology, and other fields. The examples include relations between the strength of gravity and the distance between objects, the number of earthquakes and their magnitude, the frequency of any word and its rank in the frequency table, the weight of fish and its length, and many others. Financial data also exhibit power-law regularities. Finding a theoretical explanation for power laws remains a challenge in social sciences. In our model of trading game invariance, power laws between financial variables appear naturally when time is scaled and these variables are measured in the same "units" across stocks. Similarly to physicists, we believe that the existence of power laws in financial data has a deep origin in the universal processes generating these relations and that diverse systems can be shown to share the same fundamental principles, if studied at appropriate scale.

The predictions of our model are broadly consistent with established empirical regularities in financial markets summarized in Bouchaud, Farmer, and Lillo (2008). For example, trading frequencies are found to be proportional to stock size into the power ranging between 0.44 and 0.86 (predicted to be 2/3), market impact is proportional to stock size into the power of -0.30 (predicted to be -1/3), the percentage bid-ask spread is proportional to volatility per trade (as predicted). To make these comparisons, however, we have to rely on additional assumptions that daily volatility does not vary significantly across stocks and that stock size is proportional to daily volume. Although these facts provide complimentary evidence in favor of our model, they have to be taken with a word of caution. These patterns are documented using ex post trading data - data on realized trading costs of executed trades resulted from breaking down and partial execution of unobservable orders - and therefore may be subject to systematic deviations from ex ante relations.

The remainder of this paper describes the theoretical model and empirical tests summarized above in more detail.

2 The Model

2.1 Trading Game

We develop an implementation of the continuous-time model of Kyle (1985) for the purpose of using this model to estimate from portfolio transition data how market impact varies cross-sectionally across NYSE and NASDAQ stocks with different levels of expected trading volume and expected returns volatility.

In the model of Kyle (1985), the informed trader optimally trades against noise traders and a risk-neutral market maker to exploit his private information. Trading takes place over an arbitrary period of time called a "trading day." The model delivers an intuitive benchmark for the level of equilibrium market depth. For the purpose of using this model to measure market depth empirically, however, there is no *a priori* reason to assume that this "trading day" is literally one calendar day; furthermore, the length of the trading day may vary cross-sectionally across stocks. Therefore, in our proposed model, we assume that the trading day is an endogenously determined period of time, denoted H, which might be a few seconds, a few minutes, a few hours, a few days, a few weeks, a few months, or even years. We develop an implementation of this model which is based on the intuition that the trading day H varies cross-sectionally over stocks.

The model of Kyle (1985) has two exogenous parameters: the standard deviation of fundamental value σ_V and the standard deviation of noise trading σ_U. To emphasize the dependence of these two parameters on a time period h, we shall add a subscript h to the notation and denote these parameters as $\sigma_{U,h}$ and $\sigma_{V,h}$ respectively. For $h = 1$, the notation $\sigma_{U,1}$ and $\sigma_{V,1}$ denotes standard deviations per *calendar day*, while for $h = H$, the notation $\sigma_{U,H}$ and $\sigma_{V,H}$ denote standard deviations per *trading day.*

In terms of $\sigma_{V,H}$ and $\sigma_{U,H}$, the price impact of trading x shares of stock, denoted by $\lambda \times x$, is linear, and is given by

$$\lambda = \sigma_{V,H}/\sigma_{U,H}. \tag{1}$$

Note that λ measures the price impact in dollars per share resulting from trading one share of stock; thus, λ is measured in units of dollars per share-squared. For the purpose of empirical tests and transactions cost intuition, it is useful to re-scale λ so that it is measured in basis points.

The trading activity W: In what follows, we describe how to estimate the cross-sectional variation of the parameter λ across NASDAQ and NYSE stocks with different levels of daily trading activity, which we denote as W. We define this measure as the product of the percentage daily returns volatility σ_r, the price level P, and the trading volume in shares per calendar day V:

$$W = \sigma_r \times P \times V. \tag{2}$$

According to this measure, actively traded stocks are stocks with high volatility and high dollar trading volume per calendar day, while inactively traded stocks are stocks with low volatility and low dollar trading volume per calendar day.

This measure of trading activity is consistent with the principle of Modigliani-Miller invariance, i.e. it remains unaffected by stock splits and changes in firm leverage. For example, after a two-for-one stock split, the stock price P halves but traders will trade twice as many shares, doubling V. Similarly, if the firm levers up by buying back half its outstanding shares, then volatility σ_r will double (assuming no bankruptcy) so traders will halve the quantities they trade to keep a risk per trade constant, thus halving V. In both examples, the measure of trading activity W remains the same.

In the model of Kyle (1985), the trading day measures the lifetime of private information. Our intuition is that active markets are more "efficient" than inactive markets in the sense that private information has a shorter lifetime in high volume markets and high volatility markets. In this sense, market efficiency is measured by H, with lower H representing a more efficient market. Thus, a higher level of trading activity W tends to reduce H.

The parameter $\sigma_{V,H}$: The parameter $\sigma_{V,H}$ denotes the standard deviation of private information observed by the informed trader H periods before it is revealed publicly, measured in dollars per share. Under the assumption that market makers are risk neutral, the continuous trading equilibrium has the property that prices follow Brownian motion, with the standard deviation of price changes over a trading day also equal to $\sigma_{V,H}$. The martingale property also implies that the standard deviation of price changes per calendar day, denoted $\sigma_{V,1}$, satisfies

$$\sigma_{V,H} = \sigma_{V,1} \times H^{1/2}. \tag{3}$$

The value of $\sigma_{V,1}$ can be readily estimated from data on price levels P and percentage daily returns volatility σ_r. We have

$$\sigma_{V,1} = \sigma_r P. \tag{4}$$

Note that $\sigma_{V,H}$ cannot be identified without identifying the length of the trading day H. Our intuition is that the length of the trading day H is shorter for actively traded stocks than for inactively traded stocks. As we shall see below, the length of the trading day H cannot be statistically identified from portfolio transition data. In other words, while our formulation of the model is consistent with the intuition that H declines as trading activity W increases, the parameter H remains un-identified in the econometric implementation in this paper.

The parameter $\sigma_{U,H}$: The parameter $\sigma_{U,H}$ denotes the standard deviation of the change in the inventory of noise traders measured in shares per "trading day," where noise traders are assumed to continuously place market orders so that their inventory follows a Brownian motion process. The martingale property of the inventory of noise traders implies

$$\sigma_{U,H} = \sigma_{U,1} \times H^{1/2}. \tag{5}$$

The link between the daily standard deviation of noise trading $\sigma_{U,1}$ and data on trading volume and portfolio transition trades is not straightforward because theory needs to predict how both trade frequency and trade size increase cross-sectionally with average daily volume. Our goal is to make assumptions so that $\sigma_{U,1}$ becomes identified in such a manner that it can be estimated from transition data. Even when $\sigma_{U,1}$ is identified, identification of $\sigma_{U,H}$ requires identification of H itself. The empirical tests attempt to identify $\sigma_{U,1}$ from trade sizes in portfolio transition data and daily volume data, but we do not attempt to identify $\sigma_{U,H}$ because the parameter H is not identified in our data.

Our intuition is that $\sigma_{U,h}$ is related to trading volume, but the intuition is not straightforward because the theory assumes liquidity trading follows Brownian motion but actual trades are of discrete size. The theoretical Brownian motion process for inventories implies that trading volume is infinite. For example, if we discretize trading by assuming that noise trading occurs at N discrete dates separated by time period Δt such that $N\Delta t = h$, then expected trading volume over a period of time of length h is

$$E\{\sum_{t=1}^{N} |u(t_n) - u(t_{n-1})|\} = (2Nh/\pi)^{1/2}\sigma_{U,h}. \tag{6}$$

As N becomes large, this measure of trading volume explodes.

For empirical implementation, we believe it is reasonable to approximate the Brownian motion $u(t)$ with a compound poisson process with trade arrival rate γ_1 per calendar day and distribution of trade sizes the same as some random variable denoted $\tilde{Q}$. Let $\bar{Q}$ denotes $E\{|\tilde{Q}|\}$ and let σ_Q denote the standard deviation of $\tilde{Q}$. We assume

$$\sigma_Q = \theta \bar{Q} \tag{7}$$

for some constant θ. For example, if $\tilde{Q}$ is a normal random variable, then $\theta = \sqrt{\pi/2}$. In what follows, we allow $\bar{Q}$ to vary across stocks, but we assume that θ is constant across stocks. This assumption captures the intuition that while some stocks have large average trade sizes and some stocks have small average trade sizes, the shape of the distribution of trade sizes is similar across stocks of different average trade sizes.

Over a trading day of length H, the expected number of trades γ_H is given by

$$\gamma_H = \gamma_1 \times H. \tag{8}$$

The quantity $\sigma_Q \gamma_1^{1/2}$ is the standard deviation of the change in the inventory of liquidity traders over one calendar day. The change in the inventory of liquidity traders over the trading day of length H has standard deviation

$$\sigma_{U,H} = \theta \bar{Q} \gamma_H^{1/2}, \tag{9}$$

which can equivalently be expressed as

$$\sigma_{U,H} = \theta \bar{Q} \gamma_1^{1/2} \times H^{1/2}. \tag{10}$$

TAQ Data: The assumption that the inventory of noise traders follows a Brownian motion process or a compound poisson process implies that changes in the inventory of noise traders are independently distributed. In actual trading, one independent trading decision often generates multiple reports of order executions, since trades may be broken down into smaller pieces for execution and an execution of an order may have several different counter-parties and prices.

The TAQ database gives a time-stamped record of trades printed for NYSE and NASDAQ stocks. It is probably not a good idea to estimate γ as the average number of prints in TAQ data and to estimate $\bar{Q}$ as the average print size in TAQ data. Suppose that an independent trade generates on average μ prints. Then the number of trade prints in TAQ data is $\gamma_{TAQ} = \mu\gamma$ per day, and the average trade size is $\bar{Q}_{TAQ} = \bar{Q}/\mu$. If the number of TAQ prints and the average TAQ print size are used to estimate $\bar{Q}\gamma^{1/2}$, the result is $\bar{Q}_{TAQ}\gamma_{TAQ}^{1/2} = \bar{Q}\gamma^{1/2}\mu^{-1/2}$. This estimate of $\bar{Q}\gamma^{1/2}$ is biased by a factor $\mu^{-1/2}$.

The parameter μ is not observable; moreover, it may vary across stocks. Since μ is unobservable, using average trade frequency and average trade size from TAQ data does not make it possible to calibrate the average level of price impact. If μ may vary across stocks in an unknown manner, it is not possible to use average trade frequency and average trade size from TAQ data to explain how price impact varies cross-sectionally across stocks. Whether μ is constant or varies across stocks, as a function of say stock price (based on tick size), is an interesting issue for further research.

The standard deviation of the change in the inventory of liquidity traders over one calendar day $\sigma_{U,1}$ could be also estimated from data on daily order imbalances measured as the difference between buyer initiated and seller initiated trades. Order imbalances are related to the daily trading volume but depend on its composition reflected in the number of trades, their size and direction. In theory, only a tiny fraction of trading volume is informed trading, so noise trading is almost all of observed trading volume. Thus, we expect that $\sigma_{U,1}$ can be closely approximated by the standard deviation of order imbalances. Determining order imbalances from data on trades and quotes is not straightforward because trade direction is usually unobservable. Whether empirically estimated standard deviation of order imbalances provides a reasonable alternative for estimation of market impact, is an interesting issue for future research.

2.2 Theories of Market Microstructure Invariants

The goal of our theoretical modeling is develop a framework with specific assumptions on how σ_U varies cross-sectionally across stocks with different levels of trading activity W. The distribution of trade sizes in the portfolio transition data will be used to test these assumptions given that portfolio transition trades are representative of liquidity trades implied by our theory.

The theory will then generate predictions and provide a mathematical formula for market depth and spread as functions of expected price volatility, expected average daily volume, and an unknown constant implied by the theory. The data on the implementation shortfalls in portfolio transition data will be used both to estimate the unknown constant implied by the theory and to test whether the model predicts correctly how transaction costs vary with volatility and volume.

Plugging equations (3) and (9) into equation (1) yields

$$\lambda = \frac{\sigma_r P}{\sigma_Q \gamma_1^{1/2}}. \tag{11}$$

This equation can also be written (see equation (7)) as

$$\lambda = \frac{\sigma_r P}{\theta \bar{Q} \gamma_1^{1/2}}. \tag{12}$$

We need to define several other variables before formulating our theories of invariants. Average daily volume (per calendar day), denoted V, is the product of average trade frequency γ_1 and average trade size $\bar{Q}$:

$$V = \gamma_1 \bar{Q} \tag{13}$$

Instead of operating with average trade size $\bar{Q}$, defined in number of shares and therefore affected by splits, we think of liquidity trades as bets with a given dollar standard deviation over the lifetime of the bet. This assures that liquidity trades have risk transfer properties immune to stock splits and leverage changes, thus satisfying the Modigliani-Miller invariance principle. Let liquidity "bet risk" B_1 denote the dollar standard deviation of liquidity trades. Then B_1 is given by

$$B_1 = \sigma_r P \sigma_Q. \tag{14}$$

Let B_H denote the dollar standard deviation of a liquidity trade over an entire trading day which corresponds to H caendar days. Then $B_H = B_1 \times H^{1/2}$ is given by

$$B_H = \sigma_r P \sigma_Q \times H^{1/2}. \tag{15}$$

We next describe three theories concerning how market microstructure varies across stocks. These theories are described in terms of which of the features of trading games remain invariant as the level of trading activity varies. Our proposed theory is based on the idea that the trading game itself is invariant (i.e., both the number of bets per game and their size remain constant), except for the length of time represented by the trading day over which the game is played. Our "naive" alternative theories assume either that the number of bets per calendar day are constant or that the size of liquidity traders' bets are constant.

Figure 1 Figure 1 illustrates the intuition behind the models. Let us define an arbitrary benchmark stock as a stock with daily trading volume of one million shares, price of $40 per share, and volatility of 200 basis points per day. Suppose the trading volume consists of four independent trades, each of 250,000 shares, executed over a calendar day. Each trade contributes $1/4th$ of trading volume and $1/4^{1/2}th$ of daily volatility. Therefore the size of each trade as a fraction of trading volume is equal to 1/4. The market impact of 1/4 of trading volume is equal to 100 basis points ($= 200/4^{1/2}$).

Assume now there is another stock with daily trading volume of 8 million shares, share price of $40, and daily volatility of 200 basis points. Our main question can be stated as follows. When trading volume increases by eight times (keeping other parameters constant), to what particular changes in the underlying microstructure this increase in volume corresponds? Three models give different answers on this question. To develop the intuition, we will start with our discussion of two "naive" models.

Model of Invariant Bet Frequency: Our first naive theory proposes that as average daily volume increases, average trade size $\bar{Q}$ and bet size B_1 increase proportionately but average bet frequency γ remains constant. In other words, variation in trading activity comes entirely from variation in average trade size. The assumptions imply that trade size as a share of average daily volume remains constant. Solving equation (13) for $\bar{Q}$, we find that

$$\frac{\bar{Q}_\gamma}{V} = \gamma_1^{-1} \times W^0. \tag{16}$$

To convert equation (12) into a prediction based on average daily volume and volatility, we plug the solution for $\bar{Q}$ into equation (12), obtaining

$$\lambda_\gamma = \theta^{-1} \gamma_1^{1/2} \times W^0 \times \frac{\sigma_r P}{V}. \tag{17}$$

As we shall see below, empirically there seems to be an important fixed component of trading costs, equivalent to a bid-ask spread. In the model of Kyle (1985), however, there is no explicit bid-ask spread. The discrete-time version of the model can be modified by making market makers imperfectly competitive, as in Kyle (1983). This has the effect of creating extra price impact which would not persist in a dynamic setting, capturing something like a fixed bid-ask spread. The size of this additional component of transactions is a function of the competitiveness of the market making process, as measured by the number of market makers. Since this extra component of the spread is proportional to both price impact λ and typical trade size σ_Q, we model the bid-ask spread as $\phi\lambda_\gamma\sigma_Q$, where ϕ is a constant across all stocks. The resulting solution for the bid-ask spread, denoted k, can be written

$$k_\gamma = 2\phi\gamma_1^{-1/2} \times W^0 \times \sigma_r P. \tag{18}$$

In the above equations, the subscript γ indicates that the solutions for $\bar{Q}_\gamma$, λ_γ, and k_γ hold γ_1 constant.

This naive theory is intuitively plausible. It assumes that the average trade size is a constant fraction of trading volume. It predicts that when market impact is measured in units of price standard deviation $\sigma_r P$, then the impact of trading a given percentage of average daily volume V is constant across stocks of different trading activity W. Since spread is proportional to market impact and average trade size, spread measured in units of price standard deviation $\sigma_r P$ remains constant as well.

We believe that the model of invariant bet frequency is the "default model" that implicitly but incorrectly guides the intuition of many asset managers. For example, this model justifies trading no more than 5% of average daily volume for all stocks, regardless of their level of trading activity. It also justifies attributing the same number of basis points in transactions costs for individual stocks in a basket with both active and inactive stocks, where size of trades are proportional to average daily volume.

Figure 1 illustrates the cross-sectional variation in trading games in the context of the model. The figure shows the difference between the trading game for the benchmark stock and the trading game for the stock with eight-time larger trading volume. According to the model, larger volume comes from larger trade size (keeping price and volatility constant). Similar to the benchmark stock, the stock still has four trades executed per calendar day, but their size is 2 million shares being eight times larger than trade size of the benchmark stock ($= 250,000 \times 8$). Each trade still contributes $1/4th$ of trading volume and $1/4^{1/2}th$ of daily volatility. Therefore the size of each trade as a fraction of trading volume remains equal to 1/4. The market impact of 1/4 of trading volume remains equal to 100 basis points ($= 200/4^{1/2}$).

Model of Invariant Bet Size: Our second naive theory proposes that as average daily volume increases, average trade frequency per day γ_1 increases but average bet size of horizon one day B_1 remains constant. In other words, variation in trading activity comes entirely from variation in average trade frequency. Solving equation (13) for γ_1, we find that trade size as a share of average daily volume is given by

$$\frac{\bar{Q}_B}{V} = (\theta B_1^{-1})^{-1} \times W^{-1}. \tag{19}$$

To convert equation (12) into a prediction based on average daily volume and volatility, we plug the solution for γ_1 into equation (12) and use equation (14) obtaining

$$\lambda_B = \theta^{-1}(B_1^{-1}\theta)^{1/2} \times W^{1/2} \times \frac{\sigma_r P}{V}. \tag{20}$$

Our logic from above implies that the bid-ask spread is given by

$$k_B = 2\phi(\theta B_1^{-1})^{-1/2} \times W^{-1/2} \times \sigma_r P. \tag{21}$$

In these equations, the subscript B indicates that the solutions for $\bar{Q}_B$, λ_B and k_B hold B_1 constant.

Figure 1 shows the difference between the trading game for the benchmark stock and the trading game for the stock with eight-time larger trading volume in the context of the model. According to the model, larger volume comes from higher trade frequency (keeping price and volatility constant). The stock still has trades of the same size equal to 250,000 shares as the benchmark stock. Their frequency is, however, eight times higher and equal to 32 trades per calendar day ($= 4 \times 8$). Each trade contributes $1/32th$ of trading volume and $1/32^{1/2}th$ of daily volatility. Therefore the size of each trade as a fraction of trading volume is equal to 1/32 or 1/4 times 8^{-1}. The market impact of 1/32 of trading volume is equal to $200/32^{1/2}$ basis points. This implies that the market impact of 1/4 of trading volume is equal to 100 basis points times $8^{1/2}$, calculated as $8 \times 200/32^{1/2}$ basis points. Comparing to average trade with the size equal to 1/4 of trading volume and its market impact of 100 basis points for the benchmark stock, this illustrates the appearance of -1 and $1/2$ powers in the above equations.

Both naive models make extreme assumptions attributing the variation in trading activity either entirely to the variation in bet size or entirely to the variation in bet frequency. Our proposed model of trading game invariance claims that the truth is in the middle. Neither bet size nor bet frequency remain constant as trading activity varies.

Model of Trading Game Invariance: Our proposed theory of trading game invariance assumes that both average bet frequency γ_H and average bet risk B_H are constant per trading day, not per calendar day. Intuitively, these assumptions imply that the trading game for one stock is the same as the trading game for another stock, except for the speed with which the game is played. The differences in the speed with which the game is played show up as differences in H, with small H corresponding to faster games played in more active stocks and large H corresponding to slower games played in less active stocks.

The three equations (8), (13), and (15) contain three cross-sectionally varying unobservable parameters $\bar{Q}, \gamma_1, H$, which we can solve for in terms of three observable quantities σ_r, P, V and three unobservable constants B_H, γ_H, θ. The solution expressed in terms of trading activity $W = \sigma_r P V$ is

$$H = (\gamma_H B_H \theta^{-1})^{2/3} \times W^{-2/3}, \tag{22}$$

$$\gamma_1 = (\gamma_H^{1/2} B_H^{-1} \theta)^{2/3} \times W^{2/3}, \tag{23}$$

$$\bar{Q} = (\gamma_H^{1/2} B_H^{-1} \theta)^{-2/3} \times W^{-2/3} \times V \tag{24}$$

The assumptions imply that trade size as a share of average daily volume is given by

$$\frac{\bar{Q}_{TG}}{V} = (\gamma_H^{1/2} B_H^{-1} \theta)^{-2/3} \times W^{-2/3}. \tag{25}$$

This formula is intuitive because changes in volume come from both changes in bet size and changes in bet frequency per calendar day.

Our model also implies that market depth, denoted λ_{TG} and calculated from (12), is given by

$$\lambda_{TG} = \theta^{-1}(\gamma_H^{1/2} B_H^{-1}\theta)^{1/3} \times W^{1/3} \times \frac{\sigma_r P}{V}. \tag{26}$$

When price impact is measured in units of price standard deviation $\sigma_r P$, our theoretical model predicts that the impact of trading a given percentage of average daily volume V, $\lambda_{TG} V/(\sigma_r P)$, changes across stocks different trading activity W. A one percent increase in trading activity leads to an increase of one-third of one percent in the price impact.

Similar logic for the bid-ask spread implies that the spread is given by

$$k_{TG} = 2\phi(\gamma_H^{1/2} B_H^{-1}\theta)^{-1/3} \times W^{-1/3} \times \sigma_r P. \tag{27}$$

In these equations, the subscript TG indicates that the trading game is invariant in the sense that the solutions for $\bar{Q}_{TG}$, λ_{TG}, and k_{TG} hold γ_H and B_H constant. Of course, the length of the trading day itself varies according to equation (22).

Figure 1 shows the difference between the trading game for the benchmark stock and the trading game for the stock with eight-time larger trading volume in the context of our proposed model. According to the model, when trading volume increases, the one-third of this increase comes from the increase in trade size and two-thirds of this increase comes from the increase in trade frequency (keeping price and volatility constant). Suppose trading volume increases by eight times. Trade size then increases by $8^{1/3}$ times and adds up to $500,000$ shares ($= 250,000 \times 8^{1/3}$). Trade frequency increases by $8^{2/3}$ times and amounts to 16 times ($= 4 \times 8^{2/3}$) per calendar day. Each trade contributes $1/16th$ of trading volume and $1/16^{1/2}th$ of daily volatility. Therefore the size of each trade as a fraction of trading volume is equal to $1/16$ or $1/4$ times $8^{-2/3}$. The market impact of $1/16$ of trading volume is equal to $200/16^{1/2}$ basis points or 50 basis points. This implies that the market impact of $1/4$ of trading volume is equal to $4 \times 200/16^{1/2}$ basis points or 100 basis points times $8^{1/3}$. Comparing to average trade with the size equal to $1/4$ of trading volume and its market impact of 100 basis points for the benchmark stock, this illustrates the appearance of $-2/3$ and $1/3$ powers in our model.

Model Formulation for Testing: In order to make estimated parameters have intuitive meaning, we define an arbitrary "benchmark stock". We assume that this is the same benchmark security as in Figure 1: the stock with price of \$40 per share, trading volume of one million shares per day, and volatility of 2% per day. We will also re-scale all non-identified constants.

The assumptions of three models (16), (19) and (25) can be expressed in terms of one equation that relates the trade size as a fraction of average daily volume V to the level of trading activity W,

$$\frac{\bar{Q}}{V} = \bar{q} \times \left[\frac{W}{(0.02)(40)(10^6)}\right]^{a_0}. \tag{28}$$

In this equation, the quantity $(0.02)(40)(10^6)$ in the denominator of W represents our measure of trading activity for the benchmark stock, i.e., it is the product of the 2 percent daily volatility, benchmark \$40 stock price, and one million share trading volume. Thus, the ratio of W to $(0.02)(40)(10^6)$ is one for the benchmark stock. As a result of these scaling conventions, the right-hand side is scaled so that $\bar{q}$ measures the average trade size as a fraction of daily volume for the benchmark stock.

The predictions of three models (17), (18), (20), (21), (26) and (27) can be expressed in terms of one equation that relates the implementation shortfall to the level of trading activity W. Let X denote the number of shares traded. Let $C(X)$ denote the expected cost of trading X shares of some stock, measured in basis points. We write $C(X)$ as follows:

$$C(X) = \frac{1}{2}\bar{\lambda} \times \frac{\sigma_r}{0.02}\Big[\frac{W}{(0.02)(40)(10^6)}\Big]^{\alpha_0} \times \frac{X}{(0.01)V} + \frac{1}{2}\bar{k} \times \frac{\sigma_r}{0.02}\Big[\frac{W}{(0.02)(40)(10^6)}\Big]^{\alpha_1}. \quad (29)$$

The first term on the right-hand side is the component of transactions cost due to market impact (which, if scaled to be a faction of volatility, is proportional to X given trading activity W), and the second term is the component of transactions costs due to bid-ask spread (which, if scaled to be a fraction of volatility, is constant given market activity W). The ratio of σ_r to 0.02 captures the difference in volatility of the benchmark stock and other stocks. This ratio is equal to one the benchmark stock. The ratio of W to $(0.02)(40)(10^6)$ is one for the benchmark stock as well. Similarly, the ratio of X to $(0.01)V$ is one when the trade size is one percent of average daily volume.

After re-scaling the non-identified constants, both the constant for price impact and the constant for bid-ask spread are expressed as trading costs in basis points for trading one percent of average daily volume (10,000 shares) for the "benchmark stock". We denote these constants, $\bar{\lambda}$ and $\bar{k}$, respectively. To be precise, if a trade X, representing one percent of average daily volume in the benchmark stock, incurs 8 basis points of expected costs due to price impact and 3 basis points of expected costs due to spread, then $\bar{\lambda}/2 = 8$ and $\bar{k}/2 = 3$. The total transactions cost $C(X)$ adds up to 11 basis points. Since the trade is for 10,000 shares of a \$40 stock, the 11 basis point transactions cost represents 4.4 cents per share, or \$440 for all 10,000 shares.

In defining the expected transactions cost $C(X)$, both the price impact parameter $\bar{\lambda}$ and the bid-ask spread $\bar{k}$ are divided by 2. Costs due to price impact are divided by two because the transition manager is assumed to walk up or down the demand curve, generating an average cost which is half the marginal cost represented by the price impact parameter $\bar{\lambda}$. Costs due to bid-ask spread are divided by 2 because the bid-ask spread represents a cost for a two-sided trade involving both a buy or a sell, while the one-sided trade X is either a buy or a sell, but not both.

Assumptions and predictions of all three models can be conveniently nested into the above formulation. Our proposed model of trading game invariance implies

$$\alpha_0 = 1/3, \alpha_1 = -1/3, a_0 = -2/3. \quad (30)$$

Our naive model of invariant bet frequency implies

$$\alpha_0 = 0, \alpha_1 = 0, a_0 = 0. \quad (31)$$

Our naive model of invariant bet size implies

$$\alpha_0 = 1/2, \alpha_1 = -1/2, a_0 = -1. \tag{32}$$

3 Data

3.1 Portfolio Transition Data

The empirical implications of each of the three theoretical models are tested using a proprietary database of portfolio transitions from a leading vendor of portfolio transition services. During the evaluation period, this portfolio transition vendor supervised more than 30 percent of outsourced U.S. portfolio transitions. The sample includes about 2,680 portfolio transitions executed over the period from 2001 to 2005. This database is derived from the post-transition reports prepared by transition managers for their U.S. clients. This is the same database used by Obizhaeva (2009a, 2009b).

The portfolio transitions database contains the data on individual transactions. Each observation has the following fields: a trade date, an identifier of a portfolio transition, its starting and ending dates, the name of the stock traded, the number of shares traded, buy or sell indicator, the average execution price, the pre-transition benchmark price, commissions, and fees. The data is given on separate lines for three trading venues: internal crossing networks, external crossing networks, and open market transactions. It is also given separately for each of trading days in a trading package. Old and new portfolios usually overlap. For example, both portfolios may have positions in some large and therefore widely held securities. Instead of first selling overlapping holdings from legacy portfolios and then acquiring them into target portfolios, these positions are transferred from one account to another one as "in-kind" transactions which do not incur transactions costs. Thus, if the old portfolio had 10,000 shares of IBM and the new portfolio had 4,000 shares of IBM in a portfolio transition A, then 4,000 shares are transferred in-kind and recorded as in-kind transactions. The balance of 6,000 shares will be sold. If the transition manager sells these shares in two days with open market trades on the first day and both external crosses and open market trades on the second day, then there will be 4 lines in the database corresponding to IBM stock in a given portfolio transition: a 4,000 share in-kind transaction, an open market trade the first day, an open market trade the second day, and an external cross the second day. Our empirical results do not depend at all on in-kind transfers. Instead, our empirical results are based on open market trades, external crosses, and internal crosses.

The original data is further grouped at order level. For example, aforementioned transactions are combined into one line corresponding to the order for IBM stock in portfolio transition A. This observation contains the name of the stock, the pre-transition benchmark price, buy or sell indicator, the number of shares executed over different trading venues, the average execution price for each of them, as well as the data on portfolio transition such as its beginning and ending dates.

The portfolio transition data are then matched with the CRSP to get data on stock prices, returns, and volume. Only common stocks (CRSP share codes of 10 and 11) listed on the New York Stock Exchange (NYSE), the American Stock Exchange (Amex), and NASDAQ in the period of January 2001 through December 2005 are included in the sample.

ADRs, REITS, and closed-end funds were excluded. Also excluded were stocks with missing CRSP information necessary to construct variables used for empirical tests, low-priced stocks defined as stocks with prices less than 5 dollars, and transition observations which appeared to contain typographical errors and obvious inaccuracies. Since it was unclear from the data whether adjustments for dividends and stock splits were made in a consistent manner across all transitions, all observations with non-zero payouts during the first week following the starting date of portfolio transitions were excluded from statistical tests.

After exclusions, the number of daily observations was 441,865 orders (204,780 buy orders and 237,085 sell orders).

Portfolio Transitions as Liquidity Trades The three proposed models make very different assumptions about how the sizes of liquidity trades vary across stocks with different levels of trading activity. To test these different assumptions empirically, it is necessary to identify the theoretical concept of a liquidity trade Q with actual data. Who are the liquidity traders in the stock market? One partial answer to this question is that professional equity managers are representative of liquidity traders. Although these asset managers may try to trade on the basis of private information they work hard to collect, the difficulty professional asset managers have in beating the market suggests that many of their trades do not contain much private information, and thus may be considered being liquidity trades in the context of models like Kyle (1985).

If the portfolios put together by professional asset managers result from liquidity trades, then the differences in these portfolios represent the results of numerous liquidity trades in many different stocks. Therefore, we make the identifying assumption that the differences in professionally managed portfolios, while not exactly liquidity trades themselves, vary in a manner proportional to the size of liquidity trades.

Our notation makes a distinction between the theoretical concept of a liquidity trade, denoted $\tilde{Q}$, and the individual trades made by transition managers. To emphasize the distinction, we use the notation $\bar{X}_i$ to represent the number of shares transacted in a given security during given a portfolio transition. The notation $\bar{X}_i$ represents the actual buy orders for target portfolios and the actual sell orders for legacy portfolios, excluding shares transferred in-kind. The index i ranges across 441,685 stock-transition pairs.

Portfolio transitions represent transactions in the differences between portfolios of two different professional asset managers. Note that the quantities traded often do not exactly match the positions in legacy and target portfolios. When legacy and target portfolios overlap, the overlapping positions are transferred from one account to another one as "in-kind" transactions. These in-kind transactions are transfers of positions, not trades. Therefore, these in-kind transfers are excluded from the empirical tests in this paper. As a result, the trades used in the empirical tests below represent differences in portfolio across two different asset managers. We focus on transactions rather than positions because our models are designed to explain the cross-sectional differences in the execution data. The three models establish a link between trading activity (the product of volume, price, and volatility) and trading costs with trade sizes. The models are not meant to explain the absolute levels of holdings.

Portfolio Transitions and Implementation Shortfall The fundamental problem with using implementation shortfall to measure transactions costs is that the actual quantities traded may not be known at the start date due to order cancelations or changes in trading intentions which occur after the start date and affect actual quantities traded. Statistically, the resulting selection bias problem can lead to significant underestimation of transactions costs if orders tend to be either canceled when prices move in an unfavorable direction or increased when prices move in a favorable direction. Implementation shortfall can also lead to biased estimates of transactions costs if the trading decisions are based on short-lived private information which is incorporated into prices during the period when the trades occur. Portfolio transition data has several important properties which make it particularly advantageous for estimating transactions costs using implementation shortfall.

For each stock in a portfolio transition, the quantities to be traded are known precisely at a specific time before the trades are actually executed. The composition of legacy and target portfolios is fixed in the mandates that transition managers receive the night before portfolio transitions begin. These managers then execute orders regardless of the unfolding price dynamics. This makes it reasonable to assume that the initial orders or trading intentions are exactly equal to the quantities subsequently traded. Thus, portfolio transition data tends not to be affected by the selection bias problem that would affect databases of trades where the quantities traded change in a manner correlated with price changes between the time orders are placed and the time they are executed, canceled, or increased.

The timing of portfolio transitions is likely determined by a schedule of investment committee meetings of institutional sponsors, who make decisions to undertake transitions. The investment committee meets regularly on schedules set well in advance of the meetings. Among the issues boards discuss are the replacement of fund managers and the changes of asset mix. If a decision is made to replace a portfolio manager, then a portfolio transition is arranged shortly after the meeting. These decisions are unlikely to be correlated with short-term price dynamics of individual securities during the period of the transition. This makes it possible to obtain estimates of price impact and spread that are not affected by short-lived information likely to be incorporated into prices during the period the transition trades are executed.

These properties of portfolio transitions are not often shared by other data. Consider a database built up from trades by a mutual fund, a hedge fund, or a proprietary trading desk at an investment bank. In such samples, the trading intentions of traders may not be recorded in the database. Furthermore, trading intentions before traders begin trading may not coincide with realized trades because the trader changes his mind as market conditions change. Traders often condition their trading strategies on prices by using limit orders or by canceling parts of their orders, thus hard-wiring into their strategies a selection bias problem for using such data to estimate transactions costs. The trading intentions themselves can be significantly affected by overall price dynamics, e.g., traders may be following trends or playing contrarian strategies. This dependence of actually traded quantities on prices, consequently, makes it impossible to use implementation shortfall in a meaningful way to estimate market depth and bid-ask spreads from data on trades only.

3.2 Prices, Volume, and Volatility

Our three models use trading activity to explain how transactions costs and expected trade size vary across stocks. Trading activity is the product of trading volume (in shares), share price (in dollars), and volatility (percentage standard deviation of daily returns). To measure implementation shortfall, a pre-trade benchmark price is needed. The components of trading activity and the pre-trade benchmark are calculated from CRSP data.

As a pre-trade price, denoted $P_{0,i}$, for ith trade, we use the closing price of the corresponding security on the evening before the portfolio transition trades begin. More precisely, a portfolio transition involves trades in numerous stocks. Typically, many of the stocks are traded on the first day of the transition. For each stock in the transition, the benchmark price $P_{0,i}$ is the price before the first trade is made in any of the stocks, even if a particular stock itself is not traded on the first day.

As expected trading volume during portfolio transitions, denoted V_i, for ith trade, we use the average daily trading volume (in the number of shares) of the corresponding security in the pre-transition month.

We estimate the expected volatility of daily returns, denoted $\sigma_{r,i}$, for ith trade using past daily CRSP returns for the stock involved in the ith trade. We use two different estimates of volatility, a simple estimate equal to average daily volatility from the past month and a more complicated estimate from an ARIMA model.

For each security, we first calculate the monthly standard deviation of returns from daily CRSP returns data. Let $r_{i,t,k}$ denote the CRSP return for the kth day of month t for stock involved in the ith trade. Letting $N_{i,t}$ denote the number of CRSP trading days in month t, the standard deviation for month t for stock in ith trade, denoted $\sigma^m_{i,t}$, is

$$\sigma^m_{i,t} = \big[\sum_{k=1}^{N_{i,t}} r^2_{i,t,k}\big]^{1/2} \tag{33}$$

We do not de-mean the returns data since the mean return in a month is very small relative to the standard deviation. We also do not adjust the estimates for autocorrelation of returns by adding a cross-product of adjacent returns, since this might result in the negative estimates of volatility for some stocks.

One simple estimate of daily volatility for stock in trade i for month t, denoted $\sigma^h_{i,t}$, is the monthly standard deviation converted to daily units:

$$\hat{\sigma}^h_{i,t} = \frac{1}{\sqrt{N_{i,t}}}\sigma^m_{i,t}. \tag{34}$$

We also estimate an ARIMA model to obtain another forecast of the daily return standard deviations for each stock j and month t. To reduce effects from the positive skewness of the standard deviation estimates, we use a logarithmic transformation for the volatility. We estimate a third-order moving average process for the changes in $\ln \sigma^m_{i,t}$ over the whole sample from 2001 to 2005:

$$(1-L)\ln\sigma^m_{i,t} = \Theta_0 + (1-\Theta_1 L - \Theta_2 L^2 - \Theta_3 L^3)u_t \tag{35}$$

The conditional forecast for the volatility of daily returns is

$$\hat{\sigma}^e_{i,t} = \frac{1}{\sqrt{N_{i,t}}} \exp\left[\ln \sigma^m_{i,t} + \frac{1}{2}\hat{V}(u)\right] \tag{36}$$

where $\hat{V}(u)$ is the variance of the prediction errors of the ARIMA model.

In the empirical tests below, both $\hat{\sigma}^e_{i,t-1}$ and $\hat{\sigma}^h_{i,t-1}$ are used as proxies for $\sigma_{r,i}$ in the *ith* transition trade. It is possible that using these proxies in our regressions may introduce an error-in-variables problem due to the volatility estimates themselves having errors. The empirical results are quantitatively similar for both proxies. Thus, only results for the estimates based on $\hat{\sigma}^e_{i,t-1}$ are reported. We use the pre-transition variables known before portfolio transition trades in order to avoid any spurious effects from using contemporaneous variables, except to the extent that the ARIMA model uses in-sample data to estimate model parameters.

3.3 Descriptive Statistics

Table 1 Table 1 reports statistical characteristics of both securities traded and individual transition trades. Statistics are calculated for all securities in aggregate as well as separately for ten groups of stocks sorted by average daily dollar volume. Instead of dividing the securities into ten deciles with the same number of securities, volume break points are set at the 30th, 50th, 60th, 70th, 75th, 80th, 85th, 90th, and 95th percentiles of trading volume for the universe of stocks listed on the NYSE with CRSP share codes of 10 and 11. Group 1 contains stocks in the bottom 30th percentile by dollar trading volume. Group 10 contains stocks in the top 5th percentile. Smaller percentiles for the more active stocks make it possible to focus on the stock which are most important economically. For each month, the thresholds are recalculated and the stocks are reshuffled across bins.

Panel A of Table 1 reports statistical properties of the securities in the sample. There is a column for each of the ten groupings as well as a column which reports aggregate statistics. For the entire sample of stocks, the median trading volume is \$19.99 million per day, ranging from \$1.22 million for the lowest volume decile to \$212.55 million for the highest volume decile. Since the average dollar volume ranges over more than two orders of magnitude, this variation in the data should create statistical power helpful in determining how transactions costs and trade size vary with dollar volume. Panel A reports that the median volatility for all stocks is a standard deviation in returns of 1.85 percent per day. Volatility tends to be slightly higher in the lower volume deciles than the higher ones. The volatility for the lowest volume decile is 2.04 percent, and it is 1.76 percent for the highest volume group. This implies that the amount of variation in volatility across stocks is somewhat small.

Panel A reports that the median bid-ask spread, a quoted spread obtained from the transition database, is 11.54 basis points. Its mean is 23.67 basis points. From lowest volume grouping to highest volume grouping, the median bid-ask spread declines monotonically across groups from 38.16 basis points in the lowest volume group to 4.83 basis points in the highest volume group. This monotonic decline of almost one order of magnitude in reported bid-ask spreads is so large that significant statistical power should be generated to differentiate the different predictions of the models for bid-ask spreads. This, of course, assumes that the spreads reported in Panel A, which are quoted spreads not estimated from implementation shortfall, also show up in statistical estimates based on implementation shortfall.

For example, our proposed model of trading game invariance predicts that spreads should decrease one-third of one percent for each increase of one percent in trading volume, holding volatility constant. From lowest to highest quintile, volume increases by a factor of $212.55/1.22 = 174.22$. A back-of-the envelope prediction for the decrease in spreads across these deciles groups is the one-third power of the increase in volume, i.e., $174.22^{1/3} = 5.58$. The actual decrease in spreads is a factor of $38.16/4.83 = 7.90$. While this back-of-the-envelope calculation suggests that spreads decrease more than our model predicts, the difference between 5.58 and 7.90 is small enough to warrant further statistical investigation. It is possible that the estimates of effective spreads obtained from implementation shortfall are different from quoted spreads, and thus results based on implementation shortfall will be different.

Panel B of Table 1 reports properties of order sizes in the portfolio transition data. The mean portfolio transition order is 3.90 percent of average daily volume of the stock traded. The means decline monotonically across the ten volume groups from 15.64 percent in the smallest group to 0.49 percent in the largest. The median portfolio transitions order is 0.56 percent of average daily volume. The median also declines monotonically across the ten volume groups, from 3.48 percent in the smallest to 0.14 percent in the largest. The fact that the medians are much smaller than the means indicates that the order size is skewed to the right. This is to be expected, since the order size is a non-negative number, and there may be some very small trades from highly diversified portfolios involving smaller transitions as well as very large trades from less diversified portfolios involving larger transitions.

The significant variation in mean order size as a fraction of average daily volume across dollar volume deciles is expected to have several important effects on statistical estimates.

On the one hand, the larger order size in the lower deciles generates more statistical power for using implementation shortfall to estimate market impact in the lower deciles than in the higher ones (holding constant market impact of an order with the size being a constant percentage of trading volume). On the other hand, our proposed model of trading game invariance predicts that price impact increases as volume increases, holding order size as a percentage of volume constant. Of the other two models, one predicts no change in impact while the other predicts an even larger increase. Since each of the three models makes very different predictions concerning how market impact varies with trading activity, all three models will try to extrapolate the statistical power from one volume group to another, but the extrapolation will operate differently for each model.

Second, the variation in the order size across volume groups makes it possible to test the assumptions of the three models concerning how the size of liquidity trades varies across stock with different trading activity levels. From highest to lowest group, median daily volume increases by a factor of $212.55/1.22 = 174.22$. According to our proposed model of trading game invariance, average trade size as a percent of volume should decrease by two-thirds of one percent for every one percent increase in volume, holding volatility constant. As a back-of-the envelope calculation, this implies that the decrease in trade size from lowest to highest quintile should be the two-thirds power of the increase in dollar volume, i.e., $174^{2/3} = 31.2$. The actual median trade size decreases as a fraction of average volume by a factor of $3.48/0.14 = 24.86$, from lowest group to highest group. While the back-of-the-envelope calculation of 31.2 does not exactly match the factor of 24.86, the numbers are close enough to suggest that further statistical investigation is warranted.

4 Empirical Results

All three proposed models make distinctively different assumptions concerning the cross-sectional variation in liquidity trading and offer distinctively different predictions concerning the cross-sectional variation of market impact and bid-ask spread. Portfolio transition data are used to test these implications to determine which model is more reasonable. Order size data is used to test model assumptions for liquidity order sizes. Implementation shortfall data is used to test model predictions for market impact and bid-ask spread.

4.1 Tests of Assumptions

The three theoretical models make distinctly different assumptions concerning how the size of liquidity trades varies with the level of activity. The predictions (28) can be expressed as a simple linear regression of the form

$$ln\Big[\frac{X_i}{V_i}\Big] = ln\big[\bar{q}\big] + a_0 \times ln\Big[\frac{W_i}{W_*}\Big] + \tilde{\epsilon}. \tag{37}$$

This equation relates the size of the trade X_i as a fraction of average volume V_i to the level of trading activity W_i, defined as the product of benchmark price $P_{0,i}$, last month's trading activity V_i, and estimated volatility $\sigma_{r,i}$:

$$W_i = P_{0,i} \times V_i \times \sigma_{r,i}. \tag{38}$$

The scaling constant $W_* = (40)(10^6)(0.02)$ corresponds to W_i for the hypothetical benchmark stock with price \$40 per share, trading volume of one million shares per day, and volatility of 0.02. In this regression, the model of trading game invariance predicts $a_0 = -2/3$, the model of invariant bet frequency predicts $a_0 = 0$, and the model of invariant bet size predicts $a_0 = -1$.

Table 2 Table 2 presents estimates for the coefficients in equation (37). The first column of the table reports the results of a regression pooling all the data. The four other columns in the table report results for four separate regressions in which the four parameters are estimated separately for NYSE Buys, NYSE Sells, NASDAQ Buys, and NASDAQ Sells.

The estimate for a_0 is $\hat{a}_0 = -0.63$ with standard error of 0.008 ($t = -75.27$). Economically, the point estimate for a_0 is close to the value predicted by the model of trading game invariance $a_0 = -2/3$, but this model is strongly rejected ($F = 17.03, p < 0.0001$) because the standard error is very small. This point estimate is so different from the predictions of the two other models that they are rejected by overwhelming margins.

When sample is broken down into NYSE Buys, NYSE Sells, NASDAQ Buys, and NASDAQ Sells, the point estimates -0.63, -0.60, -0.71, and -0.61 are consistently close to the predicted value of $-2/3$, but the model of trading game invariance is rejected in all cases due to the low standard error.

Table 3 Table 3 estimates the regression

$$ln\left[\frac{X_i}{V_i}\right] = ln\left[\bar{q}\right] - \frac{2}{3}ln\left[\frac{W_i}{W_*}\right] + b_1 \times ln\left[\frac{\sigma_{r,i}}{(0.02)}\right] + b_2 \times ln\left[\frac{P_i}{(40)}\right] + b_3 \times ln\left[\frac{V_i}{(10^6)}\right] + \tilde{\epsilon}. \quad (39)$$

This regression imposes on $ln(W_i/W_*)$ the coefficient $a_0 = -2/3$ predicted by the model of trading game invariance. It then allows the coefficient on the three components of W_i to vary freely. Thus, the model of trading game invariance predicts $b_1 = b_2 = b_3 = 0$. The model of invariant bet frequency predicts $b_1 = b_2 = b_3 = 2/3$, and the model of invariant bet size predicts $b_1 = b_2 = b_3 = -1/3$. The table reports point estimates for the coefficient on volatility of $\hat{b}_1 = 0.25$, the coefficient on price of $\hat{b}_2 = 0.16$, and the coefficient on share volume of $\hat{b}_3 = 0.01$, with corresponding standard errors of 0.031, 0.014, and 0.009, respectively (t-values of 8.17, 11.05, and 0.86). The regression fails to reject the hypothesis $b_3 = 0$, supporting the model of trading game invariance. But the coefficients on volatility and price are significantly positive, indicating that trade size, as a fraction of average daily volume, does not decrease with increasing volatility and volume as fast as predicted by the model of trading game invariance.

Table 4 Table 4 estimates the constant term in the regression under the assumption that the coefficient a_0 of $ln(W_i)$ in equation (37) is fixed at the values implied by the three models. For the model of trading game invariance, fixing the coefficient at $a_0 = -2/3$ results in a constant term estimate of log order size as a fraction of average daily volume, $ln\left[\bar{q}\right]$, equal to -5.69. This implies a median order size of 33.75 basis points of volume, or 0.3375% of average daily volume, for the benchmark stock.

Note that the adjusted R^2 in this one-parameter (constant term only) regression is 0.3176, compared with an adjusted R^2 of 0.3211 in the four-parameter regression and 0.3188 in the two-parameter regression. Economically, almost all the explanatory power in these regressions comes from the prediction of the model of trading game invariance that $a_0 = -2/3$.

In addition to how likely the observed data is given a particular model, it is useful to analyze which of the three models is more likely given the observed data. For each sample, the reported maximum likelihood function is greater for the model of trading game invariance than for the two alternative models. This implies that the proposed model better fits a "true" process generated empirical data according to, for example, the Akaike or Schwartz information criteria. A Bayesian statistician will assign posteriors of about 100% to the proposed model and conclude that alternative models are highly unlikely. His conclusion will be almost independent of his priors, since the extensive amount of data overwhelms priors making them irrelevant.

Figure 2 Figure 2 uses ten dummy variables for volume groups to estimate the three versions of the regression

$$\ln\left[\frac{X_i}{V_i}\right] = \left[\sum_{j=1}^{10} \mathbb{I}_{j,i} \times ln\left[\bar{q}_j\right]\right] + a_0 \times \ln\left[\frac{W_i}{W_*}\right] + \tilde{\epsilon} \quad (40)$$

in which the value of a_0 is fixed at the value predicted by the three theoretical models, and the ten dummy variables are estimated. Dummy variables $\bar{q}_j, j = 1, ..10$ quantify the average size of liquidity order for the benchmark stock based on data for jth volume group. If the regression is well-specified, then the values of the dummy variables resulting from fixing a_0 at the corresponding level should be constant across volume groups. The ten dummy variables are plotted, along with their 95% confidence bounds, on a graphs where the value of the constant term from the one-parameter regression is plotted as a horizontal line. If the regression is well-specified, then the values of the dummy variables $\bar{q}_j$ should line up along the horizontal line.

In the first graph in the figure, the ten dummy variables resulting from fixing $a_0 = -2/3$, as implied by the model of trading game invariance, are plotted. It can be seen that these dummy variables line up nicely along the horizontal line. Upon close inspection however, it is possible to notice that the 95% confidence bounds are so narrow that some of the points lie outside the 95% confidence bound, consistent with the previous rejection of the model.

In the second graph in the figure, the ten dummy variables resulting from fixing $a_0 = 0$, as implied by the model of invariant trade frequency, are plotted. Instead of lining up on the horizontal line, the dummy variables decline monotonically from a level very far above the line to a level very far below it, far outside the 95% confidence bounds. This is consistent with strong rejection of this model.

In the third graph in the figure, the ten dummy variable resulting from fixing $a_0 = -1$, as implied by the model of invariant trade frequency, are plotted. Instead of lining up on the horizontal line, the dummy variables increase monotonically from a level far below the line to a level far below it, far outside the 95% confidence bounds. This is consistent with strong rejection of this model as well.

What is the intuition underlying these patterns? For example, why the model of invariant bet frequency significantly overestimates trade size for high-volume stocks and why the model of invariant bet size significantly underestimates it? The reason is that the first "naive" model incorrectly attributes large trading volume entirely to large bet size whereas the second "naive" model mistakenly explain this volume entirely by high bet frequency. If our model of trading game invariance, however, is true, then two-thirds of this volume comes from high frequency of bets placed and one-third from their large size. Note that consistent with the two-to-one ratio between bet frequency and bet size, the deviation from a horizontal line is twice more significant for the first alternative model than for the second one.

The data on the average size of portfolio transition orders strongly support assumptions made in the model of trading game invariance and soundly reject assumptions made in alternative models. The reason is that variations in trading activity are associated with variations in both frequency and size of liquidity trades; neither remains constant.

4.2 Test of Predictions

The three theoretical models make distinctly different predictions concerning how transaction costs vary with the level of activity. These predictions can be expressed in terms of a simple equation (29) that relates the transactions cost $C(X)$ to four parameters. For a trade in the benchmark stock equal to one percent of average daily volume, the two parameters $\bar{\lambda}$ and $\bar{k}$ represent the market impact and bid-ask spread in basis points. The two remaining parameters, the exponents α_0 and α_1, describe how the models extrapolate market impact and spread costs across stocks with different levels of activity. Since the three models make dramatically different predictions concerning α_0 and α_1, it should be possible to test the models by estimating all four parameters.

We make the identifying assumption that, in a correctly specified model, the implementation shortfall from the portfolio transition database is an unbiased estimate of the transactions cost $C(X)$. We can think of implementation shortfall as representing the sum of two components: the transactions costs incurred as a result of market impact and bid-ask spread, plus the effect of random price changes between the time the benchmark price is set and the trades are executed. Since implementation shortfall is an unbiased estimate of transactions costs, we can think of the random price changes as an error in a regression. This suggests an estimation strategy of adding an error term to $C(X)$, then estimating the four parameters using a non-linear regression. The regression is non-linear because the exponent parameters α_0 and α_1 appear in $C(X)$ in a non-linear manner.

To implement this strategy, two adjustments are made, one based on statistics and one based on economics.

First, since the errors in the regression are likely to be proportional in size to the return volatility of the stock, both the right-hand-side and left-hand-side variables are divided by return volatility $\sigma_{r,i}$. This has the effect of making a crude correction for a heteroscedasticity problem which would otherwise occur. Furthermore, the imperfectly observed volatility $\sigma_{r,i}$ is replaced by its estimate $\hat{\sigma}^e_{i,t-1}$. To the extent that $\hat{\sigma}^e_{i,t-1}$ is an imperfect estimate of $\sigma_{r,i}$, the problem of bias associated with errors in variables is reduced by placing this variable on the right-hand-side.

Second, we adjust our estimation procedure to the fact that transition managers have access to different pools of liquidity: Transition orders can be executed through internal crossing networks, external crossing networks or in open markets. Market impact and bid-ask spread may be different across trading venues. Some of the portfolio transitions are the result of internal crosses. In an internal cross, one of the transition manager's customers buys from the other at some price. In fact, it is possible that both the buyer and the seller represent different portfolio transitions being implemented simultaneously. Internal crosses with other types of customers also occur, for example, crosses against flows from a passive investment management unit affiliated with the same firm as the transition management unit. Since the buyer and the seller pay the same price, it seems reasonable to assume that there is no effective spread incurred for internal crosses but there is spread for external crosses and open market transactions.

Concerning market impact for crosses and open market transactions, it is assumed that the transition manager optimally chooses the percentages of the orders to execute via these trading venues. To the extent that crosses are cheaper than open market transactions, this is expected to show up as a larger percentage of the orders being crossed than executed in open markets, not as lower market impact and spread costs on crosses. The fact that both crosses and open market transactions are used in a significant proportion of orders suggests that there are significant pools of liquidity in both crossing networks and open markets, i.e., neither dominates the other. It is thus assumed that there is market impact associated with internal crosses, of the same magnitude as with external crosses and open market trades. We will later relax the assumptions of our benchmark specification and allow transaction costs vary across trading platforms.

Table 5. Let X_i denote the number of shares in the *ith* order. Let $X_{omt,i}$ and $X_{ec,i}$ denote the number of these shares executed in open market transactions and external crosses, respectively. Then the number shares crossed internally, denoted $X_{ic,i}$ is by definition given by $X_{ic,i} = X_i - X_{omt,i} - X_{ec,i}$.

With the two aforementioned adjustments, the four parameters $\bar{\lambda}, \bar{k}, \alpha_0, \alpha_1$ are estimated in the following non-linear regression:

$$\frac{\mathbb{I}_{BS,i}(P_{ex,i} - P_{0,i})}{P_{0,i}} 10^4 \frac{(0.02)}{\sigma_{r,i}} = \frac{1}{2}\bar{\lambda} \times \left[\frac{W_i}{W_*}\right]^{\alpha_0} \frac{X_i}{(0.01)V_i} + \frac{1}{2}\bar{k} \times \frac{(X_{omt,i} + X_{ec,i})}{X_i} \left[\frac{W_i}{W_*}\right]^{\alpha_1} + \tilde{\epsilon} \quad (41)$$

In this non-linear regression, the observed data items have subscript i: $P_{ex,i}$, $P_{0,i}$, $\mathbb{I}_{BS,i}$, W_i, X_i, V_i, $\sigma_{r,i}$. Variable $\mathbb{I}_{BS,i}$ denotes trading direction being equal to 1 for buy orders and -1 for sell orders. Since $P_{0,i}$ denotes the benchmark price established the night before the transition begins and $P_{ex,i}$ denotes the average execution price, the expression $\mathbb{I}_{BS,i}(P_{ex,i} - P_{0,i})/P_{0,1} \times 10^4$ is the implementation shortfall measured in basis points. The term $(0.02)/\sigma_{r,i}$ adjusts for heteroscedasticity. The trading activity variable W_i is defined as the product of benchmark price $P_{0,i}$, last month's trading activity V_i, and estimated volatility $\sigma_{r,i}$. The scaling constant $W_* = (40)(10^6)(0.02)$ corresponds to W_i for the hypothetical benchmark stock with price \$40 per share, trading volume of one million shares per day, and volatility of 0.02. The term $X_i/(0.01)V_i$ is the size of the trade relative to average volume, scaled so that the size is one for a trade of one percent of average daily volume. The variables are scaled so that $\bar{\lambda}/2$ estimates in basis points the market impact costs of a trade of one percent of average daily volume, and $\bar{k}/2$ estimates in basis points the effective spread cost.

To adjust standard errors for positive contemporaneous correlation in returns, the 441,865 observations are pooled by week over the 2001-2005 into 4,389 clusters across 17 industry categories using the pooling option on Stata.

Recall that the three models make very different predictions concerning α_0 and α_1. The model of trading game invariance predicts $\alpha_0 = 1/3$ and $\alpha_1 = -1/3$. The model of invariant trade frequency predicts $\alpha_0 = 0$ and $\alpha_1 = 0$. The model of invariant order size predicts $\alpha_0 = 1/2$ and $\alpha_1 = -1/2$.

The results of the non-linear regression are reported in Table 5. The estimates for the parameters α_0 and α_1 are strongly supportive of the model of trading game invariance over the alternatives. The estimate for α_0 is $\hat{\alpha}_0 = 0.33$ with standard error of 0.024 ($t = 13.37$). This point estimate is almost exactly equal to the value of 1/3 predicted by the model of trading game invariance. Furthermore, the standard error is sufficiently small that predictions of the other two models, $\alpha_0 = 0.50$ and $\alpha_0 = 0$, are soundly rejected.

The estimate for α_1 is $\hat{\alpha}_1 = -0.39$ with standard error of 0.025 ($t = -15.73$). This estimate is somewhat more negative than the value $\alpha = -1/3$ predicted by the model of trading game invariance, by a margin of slightly more than two standard errors. The result suggests that effective bid-ask spreads decrease faster than the model predicts as trading activity increases. This is consistent with the back-of-the-envelope calculation from Table 1 suggesting that quoted bid-ask spreads decline faster than the model predicts as activity increases.

A Stata F-test for the joint hypothesis $\alpha_0 = 1/3$, $\alpha_1 = -1/3$ is rejected with a borderline p-value of 0.0742. Similar F-tests soundly reject the other two models with p-values less than 0.0001.

The estimate for half-price-impact is $\hat{\lambda}/2 = 2.85$ with standard error of 0.245 ($t = 11.60$), and the estimate for half-spread is $\hat{k}/2 = 6.30$ with standard error of 1.131 ($t = 6.31$). These estimates imply that a hypothetical trade in the benchmark stock equal to one percent of daily volume incurs a market impact cost of 2.85 basis points and a spread cost of 6.30 basis points. The total cost of 9.15 basis points represents 3.66 cents per share for a $40 stock, or $366 for the hypothetical 10,000 share benchmark block.

The estimate for the bid-ask spread k is double the point estimate for the half-spread $k/2$, i.e. 12.60 basis points. This estimate is somewhat higher than the median spread of 8.09 basis points reported in Table 1 for volume group 7, to which the hypothetical benchmark stock would belong. It is, however, similar to its mean value of 12.14 basis points.

Similarly, the estimate for λ is double the estimate of 2.85 for $\lambda/2$, i.e., it is 5.70. This means that a trade of 10,000 shares, one percent of average daily volume in the benchmark stock, increases the $40 price by 5.70 basis points, or 2.28 cents per share. The model implies that this increase persists over time, but it is not permanent, since the persistent effects of liquidity trades eventually dissipate due to informed trading driving the price back towards its long-term fundamental value. In the model, how fast this happens depends on the length of the trading day. In an active stock with a short trading day, markets are very resilient and the effects of noise trading are not likely to persist for long.

When the four parameters are estimated separately for NYSE Buys, NYSE Sell, NASDAQ Buys, and NASDAQ Sells, the results are also supportive of the model of trading game invariance. In three of the four regressions with the exception of NYSE Buys, the estimated coefficient for α_0 is close to the predicted value of 1/3, but α_1 is more negative than predicted. In these three cases, F-tests either fail to reject or narrowly reject the predictions of trading game invariance that $\alpha_0 = 1/3, \alpha_1 = -1/3$, with p-values of 0.1057, 0.9114, and 0.0443.

The disaggregated results for NYSE Buys, NYSE Sells, NASDAQ Buys, and NASDAQ Sells also suggest that buying is more expensive than selling. For NYSE and NASDAQ, both estimated market impact costs and estimated spread costs are larger for buy orders than for sell orders by margins that are economically meaningful if not statistically significant. For example, the effective spread for NASDAQ Buys is estimated to be more than twice as large as the effective spread for NASDAQ Sells. This is consistent with the idea that the market believes that buy orders contain more information than sell orders. See Obizhaeva (2009a) for further discussion of this idea. It is also consistent with the possibility that closing benchmark prices are biased towards the bid side of the market.

Table 6. Table 6 reports the results of a non-linear regression with a more general specification than Table 5. Three separate market impact parameters and three separate spread parameters are estimated for open market trades, external crosses, and internal crosses. In addition, the exponents on the three components of market activity (volume, price, volatility) are allowed to differ. The regression estimated is

$$
\begin{aligned}
&\frac{\mathbb{I}_{BS,i}(P_{ex,i} - P_{0,i})}{P_{0,i}} 10^4 \frac{(0.02)}{\sigma_{r,i}} \\
&\quad = \frac{1}{2} \frac{\bar{\lambda}_{omt,i} X_{omt,i} + \bar{\lambda}_{ec,i} X_{ec,i} + \bar{\lambda}_{ic,i} X_{ic,i}}{(0.01) V_i} \times \left[\frac{W_i}{W_*}\right]^{1/3} \times \frac{\sigma_{r,i}^{\beta_1} \times P_{0,i}^{\beta_2} \times V_i^{\beta_3}}{(0.02)(40)(10^6)} \\
&\quad + \frac{1}{2} \frac{\bar{k}_{omt,i} X_{omt,i} + \bar{k}_{ec,i} X_{ec,i} + \bar{k}_{ic,i} X_{ic,i}}{X_i} \times \left[\frac{W_i}{W_*}\right]^{-1/3} \times \frac{\sigma_{r,i}^{\beta_4} \times P_{0,i}^{\beta_5} \times V_i^{\beta_6}}{(0.02)(40)(10^6)} + \tilde{\epsilon}
\end{aligned} \tag{42}
$$

Because the exponents on the W-terms are set to be $1/3$ and $-1/3$, the model of trading game invariance predicts

$$\beta_1 = \beta_2 = \beta_3 = \beta_4 = \beta_5 = \beta_6 = 0. \tag{43}$$

The model of invariant bet frequency predicts

$$\beta_1 = \beta_2 = \beta_3 = -1/3, \qquad \beta_4 = \beta_5 = \beta_6 = 1/3. \tag{44}$$

The model of invariant bet size predicts

$$\beta_1 = \beta_2 = \beta_3 = 1/6, \qquad \beta_4 = \beta_5 = \beta_6 = -1/6. \tag{45}$$

The first column of the table presents the results for all buys and sells. The remaining four columns present results for separate regressions for NYSE Buys, NYSE Sells, NASDAQ Buys, and NASDAQ Sells.

F-tests of the above restrictions for the model of invariant bet frequency ($F = 78.25$) and the model of invariant bet size ($F = 13.71$) are rejected very strongly ($p << 0.0001$). An F-test of the restrictions of our model in equation (43) is rejected less strongly, with $F = 4.55$, $p = 0.0001$. From the table, it appears that one reason for this rejection is that bid-ask spreads decrease faster than predicted as trading volume increases: The estimate of β_6 is -0.09 with standard error of 0.025 ($t = 3.46$). The rapid decrease with trading volume is consistent with the results from Table 1. The bid-ask spread, however, does not decrease as fast as predicted when stock price increases, since β_5 is estimated as 0.18 with standard error of 0.062 ($t = 1.83$). Our tests do not show any significant deviation of bid-ask spread from predicted values with respect to volatility. Another reason for the rejection is that the estimates of β_1 and β_2 are quite negative. The estimate of β_1 is -0.31 with standard error of 0.191 ($t = -1.61$), and the estimate of β_2 is -0.22 with standard error of 0.191 ($t = -2.26$). The estimate of β_3 is close to zero. These results say that market impact behaves as predicted by the model of trading game invariance when the number of shares traded increases, but market impact decreases relative to what is predicted when volatility and stock price increase.

The rejection of the model of trading game invariance seems to be related to the fact that the exponents for volatility and price behave differently from the coefficients for share volume; the coefficients for volatility behave similarly to the coefficients for price. This suggests that the rejection might depend in a subtle manner on tick effects. When volatility is high and stock price is high, the tick size is small relative to a typical day's trading range.

Despite increasing the number of estimated parameters from four to twelve, the adjusted R^2 in the aggregate regression increases only from 0.0123 to 0.0129.

The estimates for the three half spread parameters are $\hat{k_{omt}}/2 = 6.56$, $\hat{k_{ec}}/2 = 6.26$, and $\hat{k_{ic}}/2 = 0.25$. These results support the assumption that there is no spread associated with internal crosses, and the spread associated with external crosses is the same as the spread associated with open market trades.

The point estimates for market impact parameters are $\hat{\lambda_{omt}}/2 = 4.49$, $\hat{\lambda_{ec}}/2 = 2.17$, and $\hat{\lambda_{ic}}/2 = 2.41$. These results support the assumption that internal crosses do have market impact. The results, however, suggest that the impact for open market trades may be greater than the impact for internal and external crosses. While interpreting these results, we have to keep in mind, however, that the separate estimates of transaction costs parameters for different trading venues may suffer from the selection bias, since transition managers optimally chooses trading venues to minimize the total costs.

Table 7. Table 7 presents estimates for equation (41) with the parameters α_0 and α_1 restricted to be as predicted in the model of trading game invariance, the model of invariant bet frequency, and the model of invariant bet size, respectively. For each of the three models, only two parameters are estimated: half-price-impact $\lambda/2$ and half spread $k/2$.

For the model of trading game invariance, the reduction from four parameters to two parameters reduces the adjusted R^2 from 0.0123 to 0.0122, consistent with very mild rejection of the model reported in Table 5. Furthermore, the parameter estimates for half-market impact $\hat{\lambda}/2$ and half bid-ask spread $\hat{k}/2$ do not change much.

For the model of invariant bet frequency, the reduction from four parameters to two parameters reduces the R^2 greatly, from 0.0123 to 0.0075, consistent with very strong rejection of the model. Furthermore, the point estimate for half market impact drops enormously, from $\hat{\lambda}/2 = 2.85$ to $\hat{\lambda}/2 = 0.3788$. This is offset by a large increase in the estimated half spread, from $\hat{k}/2 = 6.31$ to $\hat{k}/2 = 15.29$. The model of invariant bet frequency is intuitively appealing since it suggests that the market impact of trading a given percentage of average daily volume is constant as a fraction of daily returns standard deviation, regardless of the level of trading activity in the stock. The strong rejection of this model, combined with the large changes in estimated coefficients, suggests that this model leads to the misleading empirical result that market impact is less important than it really is, and bid-ask spread is more important than it really is. Therefore, one of the justifications for the model of trading game invariance is that it allows for the importance of market impact to be estimated more accurately from a better specified model.

For the model of invariant bet size, the reduction from four parameters to two parameters reduces the adjusted R^2 from 0.0123 to 0.0110, consistent with a strong rejection of this model. The point estimates for half market impact and half bid-ask spread change in the opposite direction, with the point estimate for market impact increasing from 2.85 to 3.92 and the point estimate for half spread decreasing from 6.31 to 3.46.

In all three specifications, separate regressions for NYSE Buys, NYSE Sells, NASDAQ Buys, and NASDAQ Sells continue to suggest that buying is more expensive than selling or that benchmark prices are biased towards the bid side of the market.

For each subsample, the model of trading game invariance has the largest maximum log-likelihood function among three models under consideration. This evidence implies that, according to various likelihood-based comparison methods, our model delivers predictions that are more consistent with portfolio transition data than predictions of alternative models. Almost regardless of priors, a Bayesian statistician will calculate the posterior probability of our proposed model being close to one and the posterior probabilities of alternative models are close to zero.

Figure 3 Figure 3 presents the results of three linear regressions, one for each of the three proposed models. The regression represents a modification of equation (41) in two ways. First, similarly to Table 7, for each of the three models, the values of α_0 and α_1 are fixed at the levels predicted by the models. Second, a dummy variable for each of the ten volume groups is associated with a half-market impact parameter and a half spread parameter for each group. The result is a regression with twenty coefficients, two coefficients for each volume bin, with one coefficient $\bar{\lambda}_j$ for half market impact and one coefficient $\bar{k}_j$ for half spread. The regression equation can be written

$$\frac{\mathbb{I}_{BS,i}(P_{ex,i} - P_{0,i})}{P_{0,i}} 10^4 \frac{(0.02)}{\sigma_{r,i}} = \tag{46}$$

$$\left(\sum_{j=1}^{10} \mathbb{I}_{j,i} \times {}^1\!/\!{}_2\bar{\lambda}_j\right) \times \left[\frac{W_i}{W_*}\right]^{\alpha_0} \frac{X_i}{(0.01)V_i} + \left(\sum_{j=1}^{10} \mathbb{I}_{j,i} \times {}^1\!/\!{}_2\bar{k}_j\right) \times \frac{(X_{omt,i} + X_{ec,i})}{X_i} \left[\frac{W_i}{W_*}\right]^{\alpha_1} + \tilde{\epsilon}$$

In the figure, for each of the three models, there is a graph of the estimates of the ten dummy variables for half market impact $\bar{\lambda}/2$ and a graph of the estimates of the ten dummy variables for half-spread $\bar{k}/2$. Each graph also shows the 95% confidence intervals around the point estimates, as well as a horizontal line showing the point estimate from Table 7. If the model is well specified, then the ten dummy variables should be the same and should equal the point estimates from Table 7.

For the model of trading game invariance, all of the point estimates lie either within the 95% confidence bands or slightly outside the 95% confidence bands, consistent with the mild rejection of the model discussed above. For the smallest volume group, the estimate for the half-spread has a very small confidence band, which anchors the point estimate close to the two-parameter model. For the smallest volume group, the estimate for half-market impact also has a relatively small confidence band, which anchors it close to the two-parameter model as well. For the two largest groups, the half-spread estimates are somewhat larger than the unconstrained estimate and the half-market impact estimates are somewhat smaller. The data seem to be saying that for the very largest stocks, there is a somewhat bigger spread and somewhat less market impact than implied by the model of trading game invariance. For trade groups from 2 to 6, the data are saying the opposite, i.e., that the half spread should be smaller and the half market impact larger than in the two parameter model. These patterns suggest that the transition manager might be using basket trades and then split the total costs among individual stocks in a manner consistent with the model of invariant trading frequency, i.e., assigning smaller than needed costs to large stocks and larger than needed costs to small stocks.

Similar graphs for the dummy variables in the model of trade frequency invariance are presented in the middle of the figure. This model predicts that effective bid-ask spreads do not decline as trading activity increases. It is clear from the figure, however, that the estimated effective bid-ask spreads for the smallest volume group are far greater than the estimated bid-ask spreads for the other nine groups. This places the effective spread for the smallest group far above the point estimate from the two parameter model, and very far outside the 95% confidence bands. For the market impact parameters, the model generates a great deal of power from the smallest volume group because the trade sizes are large relative to volume for this group. The point estimate of half-market impact for the smallest volume group is therefore very close to the point estimate from the two parameter model. But this forces the dummy variables for half-market impact for the nine large volume groups to lie far above the point estimate from the two parameter model. If the smallest volume group were eliminated from consideration, it appears from the figure that the model of trade frequency invariance would perform almost as well as the model of trading game invariance. It is possible that the transition manager trades baskets of stocks and then marks the prices according to the model of invariant trade frequency. Perhaps basket trades tend to occur in the largest stocks; the model might look good for spurious reasons especially among the largest stocks.

Graphs of the dummy variables for the model of invariant bet size are present on the right-hand side of the figure. The model generates very precise estimates of spreads for the smallest size group. For the larger size groups, the predicted spreads are much larger than the point estimates from the two parameter model. For half market impact, the dummy variables decrease almost monotonically, indicating that the rapid increase in market impact implied by the model ($\alpha_1 = 1/2$) is greater than what is consistent with the data.

5 Implications

Transaction Costs Formula: This paper provides a simple formula for calculation of expected transaction costs as a function of observable dollar trading volume and volatility. Using portfolio transition data, we estimate the level of transaction costs for the benchmark stock with price of \$40 per share, trading volume of one million shares per day, and volatility of 2% per day. In particular, Table 7 shows that, if the exponent parameters are set to the values implied by the model of trading game invariance, then the estimated values of half market impact $\lambda^*/2 = 2.89$ basis points and half bid-ask spread $k^*/2 = 7.91$ basis points. These estimates imply that a trade of one percent of average daily volume in the benchmark stock incurs a market impact cost of 2.89 basis points and a bid-ask spread cost of 7.91 basis points.

The model of trading game invariance then describes how transaction costs vary across stocks with different trading activity. A simple formula for expected costs $C(X)$ shows how to extrapolate the estimated transaction costs for the benchmark stock to any other security. The expected trading costs for an order of X shares, denoted $C(X)$, can be calculated as,

$$C(X) = \frac{1}{2}\lambda^* \times \left(\frac{W}{(0.02)(40)(10^6)}\right)^{1/3} \frac{\sigma_r}{0.02} \frac{X}{(0.01)V} + \frac{1}{2}k^* \times \left(\frac{W}{(0.02)(40)(10^6)}\right)^{-1/3} \frac{\sigma_r}{0.02},$$

where W is the product of share price P, daily trading volume V, and daily returns volatility σ_r. The values for $\lambda^*/2$ and $k^*/2$ are fixed at the estimated levels of 2.89 and 7.91 basis points, respectively.

We believe this formula may help to determine transaction costs not only for the U.S. stocks but also for other securities, since markets for most of them are likely to share the same underlying principles. Given the estimates $\lambda^*/2$ and $k^*/2$ for the benchmark stock, the suggested extrapolation may allow us to find expected transaction costs, for example, for fixed income and FX securities as well as for securities in other markets across the world, after adjusting our extrapolation method for differences in exchange rates. Whether the formula for transaction costs is indeed of such a general nature, is an interesting issue for future research.

Other Implications: Our paper provides answers to some important questions concerning management of transaction costs. Using the formula for expected costs and their components, asset managers can better forecast what expenses they will incur during implementation of their investment strategies as well as what amount of funds can be allocated to a strategy before it becomes unprofitable. Understanding cross-sectional variation in transaction costs has other practical implications. For example, when comparing execution quality across brokers specializing in stocks with different trading activity, performance metrics should take account of non-linearities documented in our paper. When executing basket trades, the scheme that imputes the appropriate number of basis point to individual securities in a basket has to be adjusted for cross-sectional differences in trading activity rather than assign the same number of basis points to active and inactive stocks, if size of trades is proportional to average daily volume. Finally, our analysis has implications for trading strategies: If it is reasonable to restrict trading of the benchmark stock to say 1% of average daily volume, then a smaller percentage would be appropriate for more active stocks and a larger percentage would be appropriate for less active stocks.

6 Conclusion

This paper proposes three theoretical models that differ in their assumptions about what features of trading games remain invariant as games themselves vary across securities with different levels of trading activity. Our preferred model of trading game invariance is based on the intuition that deep parameters of the trading game itself are invariant (e.g., both number of bets and their size per game), but the length of the trading games varies across stocks. Two alternative models assume that the length of these games is the same but what varies is either number of bets per game or their size. All three models make distinctive predictions concerning how market impact and bid-ask spreads vary cross-sectionally across stocks.

Data on portfolio transitions is used to test the models in two ways. First, under the identifying assumption that portfolio transitions are proportional to liquidity trades, the size of portfolio transition orders is used to test assumptions of the models concerning how liquidity order sizes varies across stock with different activity levels. Second, implications for market impact and spreads are tested using estimates derived from implementation shortfall.

The empirical results are supportive of the model of trading game invariance, but with some caveats. The model assumes that if trading activity increases by one percent, trade size as a fraction of daily volume falls by two-thirds of one percent. The trade size regressions provide strong support for this assumption. The coefficient estimate of -0.63, is remarkably close to the predicted value of -2/3. The predictions of the model for market impact and bid-ask spread are also supported by data. The empirical prediction that a one percent increase in trading activity increases the market impact (in units of daily standard deviation) by one-third of one percent is almost exactly the point estimate from non-linear regressions based on implementation shortfall. The empirical predictions that a one percent increase in trading activity decreases the bid-ask spread (in units of standard deviation) by one-third of one percent is matched by data reasonably closely.

There are, however, several issues which need further investigation. First, the statistical power behind implementation shortfall results come mostly from the 30 percent of stocks in the lowest dollar volume group. For the top 70 percent of stocks by dollar volume, it may be difficult to distinguish the model of trading game invariance from the model of invariant bet frequency. Second, in the implementation shortfall regressions, the bid-ask spreads decrease with increased activity somewhat faster than the model of trading game invariance predicts. Third, our measure of trading activity can be thought of as the product of share volume and price volatility in dollars per share. Although the model predicts that these two components of trading activity should behave similarly, both the implementation shortfall regressions and the trade size regressions suggest that they behave differently. Trading volume (measured in shares) seems to be more consistent with the model of trading game invariance than dollar price volatility. It is possible that these issues have something to do with the interaction between tick size effects and trading volume.

Interesting issues for further research include testing the three proposed models on different databases. For example, the models predictions concerning spreads can be tested using quoted spreads from TAQ data. Although it is difficult to measure the level of market depth from TAQ data using, for example, the approach of Lee and Ready (1991), the model's cross-sectional implications concerning market impact might be testable using this approach. The predictions concerning noise trading quantities can be tested using changes in holdings of mutual funds or other reporting institutional traders.

It is also possible that the model tested on stock data in this paper generalized to other markets. For example, market impact and spreads in bond markets, currency markets, or futures markets may be consistent with the regressions estimated for stocks in this paper.

References

Back, Kerry, and Shmuel Baruch, 2004, "Information in Securities Markets: Kyle meets Glosten and Milgrom," *Econometrica*, 72: 433-465.

F. Black, 1986, "Noise", Journal of Finance, vol. 41, pp. 529543 (1986)

Black, Fischer, 1995, "Equilibrium Exchanges," *Financial Analysts Journal*, May-June, 23-29.

Bouchaud, Farmer, and Lillo (2008).

Breen, Hodrick, RA Korajczyk - Management Science, 2002

Clarke, J., Shastri, K., 2001. On information asymmetry metrics. Unpublished working paper, Georgia Institute of Technology and University of Pittsburgh.

Dufour, A., and R. F. Engle. (2000). Time and the Price Impact of a Trade. Journal of Finance 55(6), 24672598.

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

Gabaix, X., Gopikrishnan, P., Plerou, V., Stanley, H. E., 2003. A theory of power law distributions in financial market fluctuations. Nature 423, 267-270.

Kyle, Albert S, 1985, "Continuous Auctions and Insider Trading", *Econometrica*, 53(6): 1315-1335.

Lin, J.-C., G. Sanger, and G. G. Booth, 1995b, Trade Size and Components of the Bid-Ask Spread, Review of Financial Studies, 8, 11531183.

Treynor, Jack (as Bagehot), 1971, "The Only Game in Town, 22: 12-14.

Hasbrouck, J., 1991, Measuring the information content of stock trades, Journal of Finance 46, 179-207.

Lee, C., and M. Ready, 1991, Inferring trade direction from intraday data, Journal of Finance 46, 733-746.

Madhavan, A., M. Richardson, and M. Roomans. "Why Do Security Prices Change? A Transaction-Level Analysis of NYSE Stocks." Review of Financial Studies, 10 (1997), 1035-64.

Obizhaeva, Anna, 2009a, Portfolio transitions and price dynamics, *University of Maryland, Working paper.*

Obizhaeva, Anna, 2009b, Selection bias in liquidity estimates, *University of Maryland, Working paper.*

Perold, A., 1988. The implementation shortfall: Paper vs. reality. Journal of Portfolio Management 14, 4-9.

Roll, R., 1984. A simple implicit measure of the effective bid-ask spread in an efficient market. Journal of Finance 39, 1127-1139.

van Ness, Bonnie, van Ness, Robert, and Warr, Richard (2001) How well do adverse selection components measure adverse selection? Financial Management 30(3), 7798.

Figure 1: Intuition Behind Models

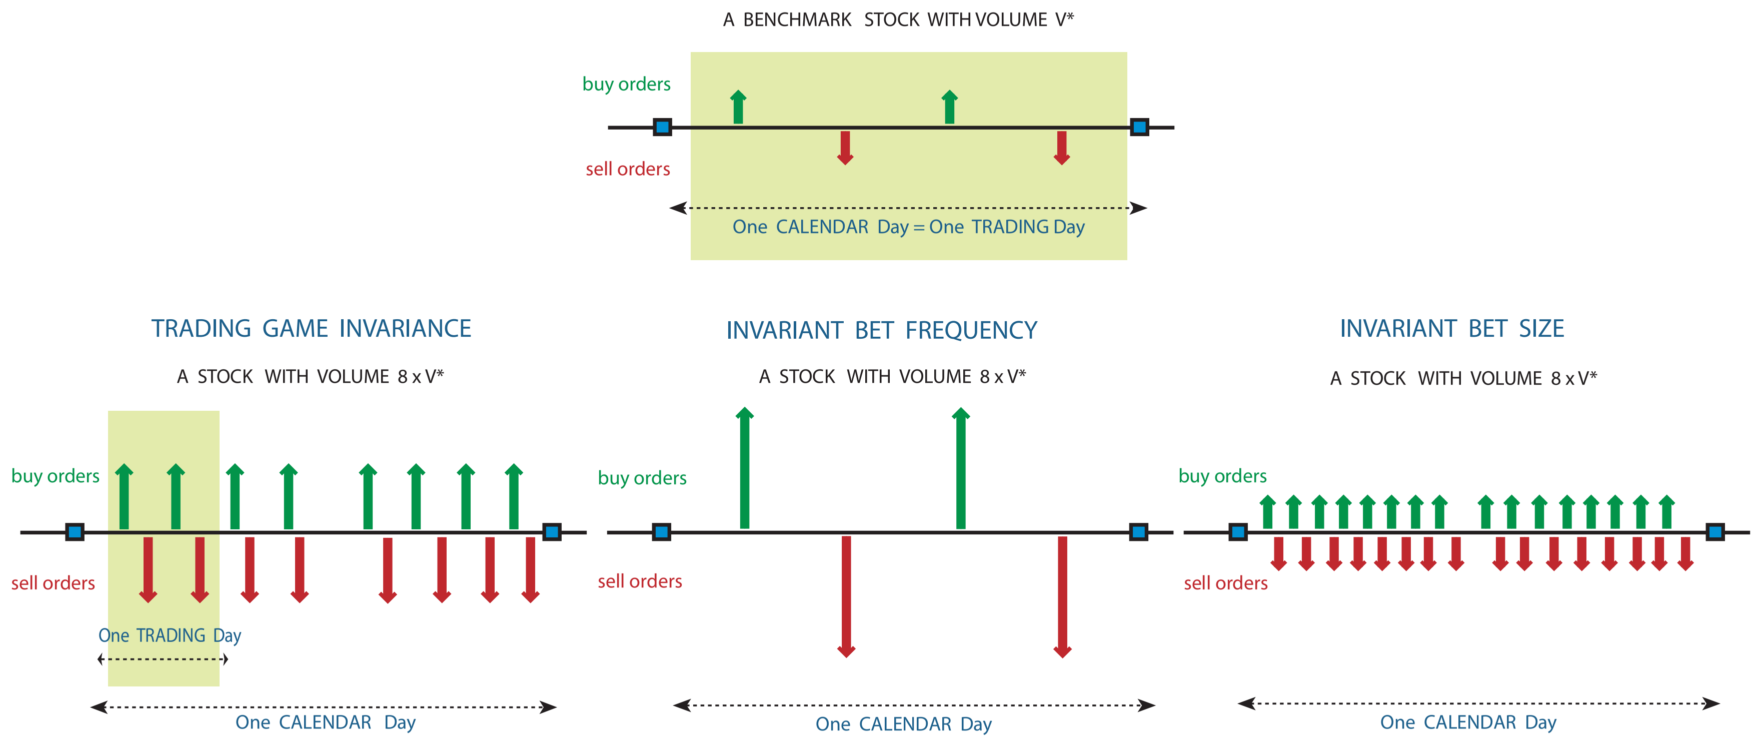

Figure illustrates the assumptions of our models. The trading game of the benchmark stock (top chart) is compared to the trading game of the stock with trading volume being eight times larger than that of the benchmark stock (bottom chart). Prices and volatilities are assumed to be the same for both stocks. The benchmark stock has four trades per calendar day; buy orders are marked in green color, and sell orders are marked in red color. The model of "invariant bet frequency" assumes that the difference in volume comes from the difference in trade sizes. The model of "invariant bet size" assumes that the difference in volume comes from the difference in trade frequencies. The model of "trading game invariance" assumes that both trade sizes and trade frequencies are the same, but the length of the game is longer for the benchmark stock.

Table 1: Descriptive Statistics

Panel A: Properties of Securities

	All	1	2	3	4	5	6	7	8	9	10
Med(V) (m $)	19.99	1.22	5.14	9.97	15.92	23.92	31.45	42.11	60.16	101.51	212.55
Med(σ_r)	1.89	2.04	2.00	1.92	1.95	1.88	1.85	1.79	1.78	1.76	1.76
Med($Sprd$) (bps)	11.54	38.16	18.34	13.53	11.81	10.12	9.34	8.09	7.16	5.92	4.83
Mean($Sprd$) (bps)	23.67	64.05	31.27	21.83	18.40	15.65	13.86	12.14	11.00	9.02	7.46

Panel B: Properties of Orders

	All	1	2	3	4	5	6	7	8	9	10
Avg(X_t/V) (%)	3.90	15.64	4.58	2.63	1.82	1.36	1.18	1.07	0.88	0.69	0.49
Med(X_t/V) (%)	0.56	3.48	1.39	0.80	0.54	0.40	0.35	0.30	0.25	0.20	0.14
Avg OMT Share	0.31	0.38	0.33	0.32	0.32	0.31	0.31	0.30	0.29	0.27	0.24
Avg EC Share	0.40	0.42	0.42	0.41	0.41	0.41	0.40	0.40	0.39	0.38	0.36
Avg IC Share	0.29	0.20	0.25	0.26	0.27	0.28	0.29	0.30	0.32	0.35	0.40
# Obs	441,865	65,081	68,545	41,559	49,532	28,621	30,087	30,710	35,733	42,331	49,666

Table reports the characteristics of securities and transition orders in the sample. Panel A shows the median of average daily dollar volume V (in millions of $), the median of the daily returns volatility σ_r in percents, the median and the mean of the percentage spread $Sprd$ in basis points. Panel B shows the average order size (in percents of V), the median order size (in percents of V), the average fraction of transition order executed in open market (Avg OMT Share), external and internal crossing networks (Avg EC and IC Shares), as well as the total number of observations. Results are presented for stocks with different dollar trading volume. Group 1 (Group 10) contains orders in stocks with lowest (highest) dollar trading volume. Each month, the observations are split into 10 bins according to stocks' dollar trading volume in pre-transition month. The thresholds correspond to 30th, 50th, 60th, 70th, 75th, 80th, 85th, 90th, and 95th percentiles of dollar trading volume for common stocks listed on the NYSE. The sample ranges from January 2001 to December 2005.

Table 2: Model Testing for Order Sizes I

		NYSE		NASDAQ	
	All	Buy	Sell	Buy	Sell
$ln\left[\bar{q}\right]$	-5.67***	-5.68***	-5.63***	-5.75***	-5.65***
	(0.017)	(0.022)	(0.018)	(0.033)	(0.031)
a_0	-0.63***	-0.63***	-0.60***	-0.71***	-0.61***
	(0.008)	(0.010)	(0.008)	(0.019)	(0.012)
	Model of Trading Game Invariance: $a_0 = -2/3$				
F-test	17.01	13.74	72.00	6.53	18.56
p-val	0.0000	0.0002	0.0000	0.0107	0.0000
	Model of Invariant Bet Frequency: $a_0 = 0$				
F-test	5664.91	3740.45	5667.60	1440.32	2427.51
p-val	0.0000	0.0000	0.0000	0.0000	0.0000
	Model of Invariant Bet Size: $a_0 = -1$				
F-test	1920.13	1306.11	2537.08	229.30	966.99
p-val	0.0000	0.0000	0.0000	0.0000	0.0000
Q_*/V_*	34.68	34.08	35.96	31.85	35.27
$d/g/n$	2/1/4389	2/1/4018	2/1/4198	2/1/2855	2/1/2977
#Obs	441,865	135,006	152,701	69,774	84,384
Adj. R^2	0.3188	0.2588	0.2643	0.4364	0.3648
R^2	0.3188	0.2588	0.2643	0.4364	0.3648

Table presents the estimates $\bar{q}, a_0$ for the regression $ln\left[Y_i\right] = ln\left[\bar{q}\right] + a_0 \times ln\left[\frac{W_i}{W_*}\right] + \tilde{\epsilon}$. Each observation corresponds to order i. Y_i is the size of the trade relative to expected daily volume calculated as $X_i/(0.01)V_i$ where expected daily volume V_i is measured as the last month's average daily volume, the size is one for a trade of one percent of expected volume. The trading activity W_i is the product of expected daily volatility $\sigma_{r,i}$, benchmark price $P_{0,i}$, and expected daily volume V_i. The scaling constant $W_* = (0.02)(40)(10^6)$ corresponds to W_i for the benchmark stock with volatility of 0.02, price \$40 per share, and trading volume of one million shares per day. $\bar{q}$ is the measure of order size such that the median order size Q_*/V_* for a benchmark stock is calculated as $\bar{q} \times 10^4$ in basis points. The standard errors are clustered at weekly levels for 17 industries and shown in parentheses. F-statistics and p-values are reported for three models with d parameters, g restrictions, and n clusters in the regression. The sample ranges from January 2001 to December 2005. ***, **, *denotes significance at 1%, 5% and 10% levels, respectively.

Table 3: Model Testing for Order Sizes II

		NYSE		NASDAQ	
	All	Buy	Sell	Buy	Sell
$ln[\bar{q}]$	-5.65***	-5.66***	-5.58***	-5.80***	-5.63***
	(0.016)	(0.022)	(0.018)	(0.034)	(0.029)
b_1	0.25***	0.30***	0.36**	0.17*	0.24**
	(0.031)	(0.041)	(0.036)	(0.073)	(0.073)
b_2	0.16***	0.11***	0.21***	0.02	0.22***
	(0.014)	(0.018)	(0.014)	(0.039)	(0.028)
b_3	0.01	0.02	0.03**	-0.07***	0.02
	(0.009)	(0.011)	(0.009)	(0.018)	(0.015)
	Model of Trading Game Invariance: $b_1 = b_2 = b_3 = 0$				
F-test	47.57	19.92	79.29	8.77	20.97
p-val	0.0000	0.0000	0.0000	0.0000	0.0000
	Model of Invariant Bet Frequency: $b_1 = b_2 = b_3 = 2/3$				
F-test	2044.60	1286.85	1939.03	567.85	808.05
p-val	0.0000	0.0000	0.0000	0.0000	0.0000
	Model of Invariant Bet Size: $b_1 = b_2 = b_3 = -1/3$				
F-test	747.74	465.71	947.01	78.19	325.63
p-val	0.0000	0.0000	0.0000	0.0000	0.0000
Q_*/V_*	35.31	34.98	32.99	30.35	35.95
$d/g/n$	4/3/4389	4/3/4018	4/3/4198	4/3/2855	4/3/2977
#Obs	441,865	135,006	152,701	69,774	84,384
Adj. R^2	0.3211	0.2614	0.2682	0.4382	0.3674
R^2	0.3213	0.2616	0.2684	0.4384	0.3676

Table presents the estimates $\bar{q}, b_1, b_2, b_3$ for the regression $ln\Big[Y_i\Big] = ln\big[\bar{q}\big] - \frac{2}{3} \times ln\Big[\frac{W_i}{W_*}\Big] + b_1 \times ln\Big[\frac{\sigma_i^e}{(0.02)}\Big] + b_2 \times ln\Big[\frac{P_i}{(40)}\Big] + b_3 \times ln\Big[\frac{V_i}{(10^6)}\Big] + \tilde{\epsilon}$. Each observation corresponds to order i. Y_i is the size of the trade relative to expected daily volume calculated as $X_i/(0.01)V_i$ where expected daily volume V_i is measured as the last month's average daily volume, the size is one for a trade of one percent of expected volume. The trading activity W_i is the product of expected daily volatility $\sigma_{r,i}$, benchmark price $P_{0,i}$, and expected daily volume V_i. The scaling constant $W_* = (0.02)(40)(10^6)$ corresponds to W_i for the benchmark stock with volatility of 0.02, price \$40 per share, and trading volume of one million shares per day. $\bar{q}$ is the measure of order size such that the median order size Q_*/V_* for a benchmark stock is calculated as $\bar{q} \times 10^4$ in basis points. The standard errors are clustered at weekly levels for 17 industries and shown in parentheses. F-statistics and p-values are reported for three models with d parameters, g restrictions, and n clusters in the regression. The sample ranges from January 2001 to December 2005. ***, **, *denotes significance at 1%, 5% and 10% levels, respectively.

Table 4: Model Calibration for Order Sizes

	All	NYSE Buy	NYSE Sell	NASDAQ Buy	NASDAQ Sell
		Model of Trading Game Invariance: $a_0 = -2/3$			
$ln\left[\bar{q}\right]$	-5.69***	-5.70***	-5.67***	-5.70***	-5.71***
	(0.018)	(0.022)	(0.019)	(0.039)	(0.035)
Q_*/V_*	33.75	33.35	34.61	33.60	32.99
Adj. R^2	0.3177	0.2577	0.2608	0.4343	0.3618
R^2	0.3178	0.2579	0.2609	0.4345	0.3620
$Log(L)$	-828,757	-255,637	-287,149	-127,591	-158,149
		Model of Invariant Bet Frequency: $a_0 = 0$			
$ln\left[\bar{q}\right]$	-5.17***	-5.33***	-5.29***	-4.95***	-4.88***
	(0.021)	(0.025)	(0.021)	(0.047)	(0.043)
Q_*/V_*	56.85	48.44	50.42	70.83	75.97
Adj. R^2	-0.0002	-0.0002	-0.0002	-0.0004	-0.0003
R^2	0.0000	0.0000	0.0000	0.0000	0.0000
$Log(L)$	-913,255	-275,771	-310,235	-147,479	-177,111
		Model of Invariant Bet Size: $a_0 = -1$			
$ln\left[\bar{q}\right]$	-5.95***	-5.89***	-5.86***	-6.07***	-6.13***
	(0.021)	(0.024)	(0.022)	(0.045)	(0.045)
Q_*/V_*	26.05	27.67	28.51	23.11	21.77
Adj. R^2	0.2105	0.1683	0.1458	0.3669	0.2192
R^2	0.2107	0.1685	0.1460	0.3671	0.2195
$Log(L)$	-860,973	-263,318	-298,186	-131,520	-166,656
#Obs	441,865	135,006	152,701	69,774	84,384

Table presents the estimates $ln\left[\bar{q}\right]$ for the regression $ln\left[Y_i\right] = ln\left[\bar{q}\right] + a_0 \times ln\left[\frac{W_i}{W_*}\right] + \tilde{\epsilon}$ with a_0 restricted to be as predicted in proposed models. Each observation corresponds to order i. Y_i is the size of the trade relative to expected daily volume calculated as $X_i/(0.01)V_i$ where expected daily volume V_i is measured as the last month's average daily volume, the size is one for a trade of one percent of expected volume. The trading activity W_i is the product of expected daily volatility $\sigma_{r,i}$, benchmark price $P_{0,i}$, and expected daily volume V_i. The scaling constant $W_* = (0.02)(40)(10^6)$ corresponds to W_i for the benchmark stock with volatility of 0.02, price \$40 per share, and trading volume of one million shares per day. $\bar{q}$ is the measure of order size such that the median order size Q_*/V_* for a benchmark stock is calculated as $\bar{q} \times 10^4$ in basis points. The standard errors are clustered at weekly levels for 17 industries and shown in parentheses. The log-likelihood and R^2 are shown for each model. The sample ranges from January 2001 to December 2005. ***, **, *denotes significance at 1%, 5% and 10% levels, respectively.

Table 5: Model Testing for Market Impact and Bid-Ask Spread I

		NYSE		NASDAQ	
	All	Buy	Sell	Buy	Sell
$1/2\bar{\lambda}$	2.85***	2.50***	2.33***	4.20***	2.99***
	(0.245)	(0.515)	(0.365)	(0.753)	(0.662)
α_0	0.33***	0.18***	0.33***	0.33***	0.35***
	(0.024)	(0.045)	(0.054)	(0.053)	(0.045)
$1/2\bar{k}$	6.30***	14.94***	2.82*	8.38*	3.94**
	(1.131)	(2.529)	(1.394)	(3.328)	(1.498)
α_1	-0.39***	-0.19***	-0.46***	-0.36***	-0.45***
	(0.025)	(0.045)	(0.061)	(0.061)	(0.047)
	Model of Trading Game Invariance: $\alpha_0 = 1/3, \alpha_1 = -1/3$				
F-test	2.62	8.51	2.25	0.09	3.12
p-val	0.0731	0.0002	0.1057	0.9114	0.0443
	Model of Invariant Bet Frequency: $\alpha_0 = 0, \alpha_1 = 0$				
F-test	176.14	14.79	47.03	33.11	71.06
p-val	0.0000	0.0000	0.0000	0.0000	0.0000
	Model of Invariant Bet Size: $\alpha_0 = 1/2, \alpha_1 = -1/2$				
F-test	30.30	39.63	5.23	7.21	5.92
p-val	0.0000	0.0000	0.0054	0.0007	0.0027
$d/g/n$	4/2/4389	4/2/4018	4/2/4198	4/2/2855	4/2/2977
#Obs	441,865	135,006	152,701	69,774	84,384
R^2	0.0126	0.0136	0.0067	0.0211	0.0195
Adj. R^2	0.0123	0.0134	0.0064	0.0208	0.0192

Table presents the estimates for $\bar{\lambda}, \bar{k}, \alpha_0, \alpha_1$ in the regression $Y_i = \frac{1}{2}\bar{\lambda} \times \left[\frac{W_i}{W_*}\right]^{\alpha_0} \frac{X_i}{(0.01)V_i} + \frac{1}{2}\bar{k} \times \frac{(X_{omt,i}+X_{ec,i})}{X_i}\left[\frac{W_i}{W_*}\right]^{\alpha_1} + \tilde{\epsilon}$. Each observation corresponds to order i. Y_i is the implementation shortfall in basis points calculated as $\frac{\mathbb{I}_{BS,i}(P_{ex,i}-P_{0,i})}{P_{0,i}} 10^4 \frac{(0.02)}{\sigma_{r,i}}$, where $\mathbb{I}_{BS,i}$ is a buy/sell indicator, $P_{ex,i}$ is the average execution price, $P_{0,i}$ is the pre-transition price, $\sigma_{r,i}$ is the expected daily volatility estimated as $\sigma^e_{i,t-1}$. The term $(0.02)/\sigma_{r,i}$ adjusts for heteroscedasticity. The trading activity W_i is the product of expected volatility $\sigma_{r,i}$, benchmark price $P_{0,i}$, and expected volume V_i measured as last month's average daily volume. The scaling constant $W_* = (0.02)(40)(10^6)$ corresponds to W_i for the benchmark stock with volatility of 0.02, price $40 per share, and trading volume of one million shares per day. X_i is the number of shares in the order with $X_{omt,i}$ executed in open market and $X_{ec,i}$ executed in external crossing networks. The term $X_i/(0.01)V_i$ is the size of the trade relative to average volume, the size is one for a trade of one percent of expected daily volume. $\bar{\lambda}/2$ estimates in basis points the market impact costs of a trade of one percent of average daily volume in a benchmark stock, and $\bar{k}/2$ estimates in basis points the effective spread cost. The standard errors are clustered at weekly levels for 17 industries and shown in parentheses. F-statistics and p-values are reported for three models with d parameters, g restrictions, and n clusters in the regression. The sample ranges from January 2001 to December 2005. ***, **, *denotes significance at 1%, 5% and 10% levels, respectively.

Table 6: Model Testing for Market Impact and Bid-Ask Spread II

		NYSE		NASDAQ	
	All	Buy	Sell	Buy	Sell
$^1\!/_2\bar{\lambda}_{omt}$	4.49***	2.41***	4.35***	5.39***	4.33***
	(0.601)	(0.433)	(0.798)	(1.449)	(1.058)
$^1\!/_2\bar{\lambda}_{ec}$	2.17***	3.02***	1.77**	3.64***	1.34
	(0.386)	(0.471)	(0.542)	(0.968)	(0.705)
$^1\!/_2\bar{\lambda}_{ic}$	2.40***	2.20***	1.77***	2.07*	1.51**
	(0.334)	(0.622)	(0.420)	(0.891)	(0.504)
β_1	-0.31	-0.86***	-0.37	-0.10	-1.05***
	(0.191)	(0.141)	(0.327)	(0.347)	(0.284)
β_2	-0.22*	-0.01	-0.43*	-0.17	-0.32
	(0.191)	(0.141)	(0.327)	(0.347)	(0.284)
β_3	0.04	-0.19***	0.13	0.01	-0.04
	(0.043)	(0.031)	(0.077)	(0.053)	(0.037)
$^1\!/_2\bar{k}_{omt}$	6.56***	18.55***	3.05*	14.43***	4.69*
	(1.227)	(3.539)	(1.279)	(3.967)	(1.819)
$^1\!/_2\bar{k}_{ec}$	6.26***	8.99***	4.98**	11.13**	5.08**
	(1.124)	(2.355)	(1.520)	(3.690)	(1.841)
$^1\!/_2\bar{k}_{ic}$	0.26	5.31	-4.38**	7.78	0.70
	(1.846)	(3.987)	(1.429)	(7.699)	(1.317)
β_4	0.10	-0.06	0.60*	-0.29	0.99***
	(0.160)	(0.249)	(0.260)	(0.297)	(0.258)
β_5	0.18**	-0.22	0.06	0.26	0.36***
	(0.062)	(0.172)	(0.123)	(0.139)	(0.104)
β_6	-0.09***	0.26***	-0.12*	0.05	-0.11*
	(0.025)	(0.055)	(0.054)	(0.050)	(0.049)
	Model of Trading Game Invariance: $\beta_1 = \beta_2 = \beta_3 = \beta_4 = \beta_5 = \beta_6 = 0.$				
F-test	4.59	20.96	2.80	1.38	11.26
p-val	0.0001	0.0000	0.0102	0.2168	0.0000
	Model of Invariant Bet Frequency: $\beta_1 = \beta_2 = \beta_3 = -1/3, \beta_4 = \beta_5 = \beta_6 = 1/3.$				
F-test	78.26	12.06	26.46	10.71	23.27
p-val	0.0000	0.0000	0.0000	0.0000	0.0000
	Model of Invariant Bet Size: $\beta_1 = \beta_2 = \beta_3 = 1/6, \beta_4 = \beta_5 = \beta_6 = -1/6.$				
F-test	13.77	44.25	5.94	6.85	28.79
p-val	0.0000	0.0000	0.0000	0.0000	0.0000
$d/g/n$	12/6/4389	12/6/4018	12/6/4198	12/6/2855	12/6/2977
#Obs	441,865	135,006	152,701	69,774	84,384
Adj. R^2	0.0129	0.0147	0.0076	0.0222	0.0214
R^2	0.0131	0.0150	0.0079	0.0225	0.0217

Table presents the estimates for $\bar{\lambda}_{omt}, \bar{\lambda}_{ec}, \bar{\lambda}_{ic}, \bar{k}_{omt}, \bar{k}_{ec}, \bar{k}_{ic}, \beta_1, \beta_2, \beta_3, \beta_4, \beta_5, \beta_6$ in the regression $Y_i = \frac{1}{2}\frac{\lambda_{omt,i}X_{omt,i}+\lambda_{ec,i}X_{ec,i}+\lambda_{ic,i}X_{ic,i}}{(0.01)V_i}\left[\frac{W_i}{W_*}\right]^{1/3}\frac{\sigma_i^{\beta_1}P_{0,i}^{\beta_2}V_i^{\beta_3}}{(0.02)(40)(10^6)} + \frac{1}{2}\frac{k_{omt,i}X_{omt,i}+k_{ec,i}X_{ec,i}+k_{ic,i}X_{ic,i}}{X_i}\left[\frac{W_i}{W_*}\right]^{-1/3}\frac{\sigma_i^{\beta_4}P_{0,i}^{\beta_5}V_i^{\beta_6}}{(0.02)(40)(10^6)} + \tilde{\epsilon}$. Each observation corresponds to order i. Y_i is the implementation shortfall in basis points calculated as $\frac{\mathbb{I}_{BS,i}(P_{ex,i}-P_{0,i})}{P_{0,i}}10^4\frac{(0.02)}{\sigma_{r,i}}$, where $\mathbb{I}_{BS,i}$ is a buy/sell indicator, $P_{ex,i}$ is the average execution price, $P_{0,i}$ is the pre-transition price, $\sigma_{r,i}$ is the expected daily volatility estimated as $\sigma^e_{i,t-1}$. The term $(0.02)/\sigma_{r,i}$ adjusts for heteroscedasticity. The trading activity W_i is the product of expected volatility $\sigma_{r,i}$, benchmark price $P_{0,i}$, and expected volume V_i measured as last month's average daily volume. The scaling constant $W_* = (0.02)(40)(10^6)$ corresponds to W_i for the benchmark stock with volatility of 0.02, price \$40 per share, and trading volume of one million shares per day. X_i is the number of shares in the order with $X_{omt,i}$ executed in open market, $X_{ec,i}$ executed in external crossing networks, $X_{ic,i}$ executed in internal crossing networks. The term $X_i/(0.01)V_i$ is the size of the trade relative to expected daily volume, the size is one for a trade of one percent of expected daily volume. $\bar{\lambda}_{omt}/2, \bar{\lambda}_{ec}/2, \bar{\lambda}_{ic}/2$ estimate in basis points the market impact costs of a trade of one percent of average daily volume in a benchmark stock for open market trades, external crosses, and internal crosses. $\bar{k}_{omt}/2, \bar{k}_{ec}/2, \bar{k}_{ic}/2$ estimate in basis points the effective spread cost for open market trades, external crosses, and internal crosses. The standard errors are clustered at weekly levels for 17 industries and shown in parentheses. F-statistics and p-values are reported for three models with d parameters, g restrictions, and n clusters in the regression. The sample ranges from January 2001 to December 2005. ***, **, *denotes significance at 1%, 5% and 10% levels, respectively.

Table 7: Model Calibration for Market Impact and Bid-Ask Spread

		NYSE		NASDAQ	
	All	Buy	Sell	Buy	Sell
	Model of Trading Game Invariance: $\alpha_0 = 1/3, \alpha_1 = -1/3$				
$1/2\hat{\lambda}$	2.8898***	3.4199***	2.3411***	4.2361***	2.7800***
	(0.1948)	(0.4421)	(0.3231)	(0.4342)	(0.3300)
$1/2\hat{k}$	7.9036***	10.9695***	4.7553***	9.2870***	7.3187***
	(0.6889)	(1.3724)	(1.1065)	(1.5153)	(1.0067)
Adj. R^2	0.0122	0.0129	0.0062	0.0207	0.0184
R^2	0.0125	0.0131	0.0064	0.0211	0.0188
$Log(L)$	1061,808	324,590	366,995	167,793	202,588
	Model of Invariant Bet Frequency: $\alpha_0 = 0, \alpha_1 = 0$				
$1/2\hat{\lambda}$	0.3788***	1.3291***	0.4712***	1.1857***	0.2694***
	(0.0884)	(0.1130)	(0.1041)	(0.1377)	(0.0457)
$1/2\hat{k}$	15.2686***	19.1143***	6.4113**	19.2409***	13.1085***
	(1.5658)	(2.5247)	(2.3157)	(4.4536)	(3.5423)
Adj. R^2	0.0075	0.0126	0.0038	0.0178	0.0105
R^2	0.0077	0.0128	0.0040	0.0182	0.0108
$Log(L)$	1060,752	324,568	366,811	167,689	202,247
	Model of Invariant Bet Size: $\alpha_0 = 1/2, \alpha_1 = -1/2$				
$1/2\hat{\lambda}$	3.9202***	3.9578***	2.8730***	6.5057***	4.6229***
	(0.3036)	(0.6274)	(0.4339)	(0.6467)	(0.6037)
$1/2\hat{k}$	3.4648***	5.7598***	2.5020***	4.2232***	3.0040***
	(0.2917)	(0.6487)	(0.4105)	(0.5461)	(0.3587)
Adj. R^2	0.0110	0.0108	0.0056	0.0195	0.0182
R^2	0.0112	0.0110	0.0058	0.0198	0.0185
$Log(L)$	1061,533	324,446	366,950	167,748	202,578
#Obs	441,865	135,006	152,701	69,774	84,384

Table presents the estimates for $\bar{\lambda}, \bar{k}$ in the regression $Y_i = \frac{1}{2}\bar{\lambda} \times \left[\frac{W_i}{W_*}\right]^{\alpha_0} \frac{X_i}{(0.01)V_i} + \frac{1}{2}\bar{k} \times \frac{(X_{omt,i}+X_{ec,i})}{X_i}\left[\frac{W_i}{W_*}\right]^{\alpha_1} + \tilde{\epsilon}$ with α_0 and α_1 restricted to be as predicted in proposed models. Each observation corresponds to order i. Y_i is the implementation shortfall in basis points calculated as $\frac{\mathbb{I}_{BS,i}(P_{ex,i}-P_{0,i})}{P_{0,i}} 10^4 \frac{(0.02)}{\sigma_{r,i}}$, where $\mathbb{I}_{BS,i}$ is a buy/sell indicator, $P_{ex,i}$ is the average execution price, $P_{0,i}$ is the pre-transition price, $\sigma_{r,i}$ is the expected daily volatility estimated as $\sigma^e_{i,t-1}$. The term $(0.02)/\sigma_{r,i}$ adjusts for heteroscedasticity. The trading activity W_i is the product of expected volatility $\sigma_{r,i}$, benchmark price $P_{0,i}$, and expected volume V_i measured as last month's average daily volume. The scaling constant $W_* = (0.02)(40)(10^6)$ corresponds to W_i for the benchmark stock with volatility of 0.02, price \$40 per share, and trading volume of one million shares per day. X_i is the number of shares in the order with $X_{omt,i}$ executed in open market and $X_{ec,i}$ executed in external crossing networks. The term $X_i/(0.01)V_i$ is the size of the trade relative to expected volume, the size is one for a trade of one percent of expected daily volume. $\bar{\lambda}/2$ estimates in basis points the market impact costs of a trade of one percent of average daily volume in a benchmark stock, and $\bar{k}/2$ estimates in basis points the effective spread cost. The standard errors are clustered at weekly levels for 17 industries and shown in parentheses. The log-likelihood and R^2 are shown for each model. The sample ranges from January 2001 to December 2005. ***, **, *denotes significance at 1%, 5% and 10% levels, respectively.

Figure 2: Order Size across 10 Volume Groups

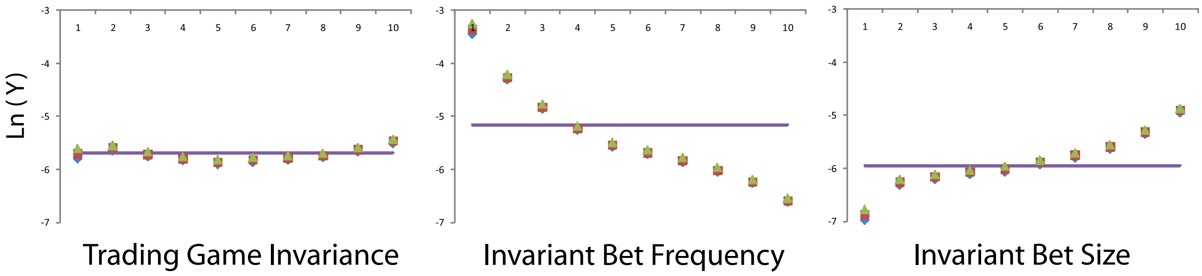

Figure shows the average logarithm of the order sizes $ln[\bar{q}_j]$ with the 95%-confidence intervals for 10 volume groups for Model 1, Model 2, and Model 3 from regression $\ln\Big[Y_i\Big] = \Big[\sum_{j=1}^{10} \mathbb{I}_{j,i} \times ln[\bar{q}_j]\Big] + a_0 \times \ln\Big[\frac{W_i}{W_*}\Big] + \tilde{\epsilon}$ as well as unconditional estimates from Table 4. In Model of Trading Game Invariance, $\alpha_0 = -2/3$. In Model of Invariant Bet Frequency, $\alpha_0 = 0$. In Model of Invariant Bet Size, $\alpha_0 = -1$. Each observation corresponds to order i. $\mathbb{I}_{j,i}$ is an indicator equal to one if order i is executed in a stock from group j. Y_i is the size of the trade relative to expected daily volume calculated as $X_i/(0.01)V_i$ where expected daily volume V_i is measured as the last month's average daily volume, the size is one for a trade of one percent of expected volume. The trading activity W_i is the product of expected daily volatility $\sigma_{r,i}$, benchmark price $P_{0,i}$, and expected daily volume V_i. The scaling constant $W_* = (0.02)(40)(10^6)$ corresponds to W_i for the benchmark stock with volatility of 0.02, price \$40 per share, and trading volume of one million shares per day. $\bar{q}_j$ is the measure of order size such that the median order size Q_*/V_* for a benchmark stock is calculated as $\bar{q}_j \times 10^4$ in basis points for volume group j. Volume groups are based on the pre-transition dollar trading volume with thresholds corresponding to 30th, 50th, 60th, 70th, 75th, 80th, 85th, 90th, and 95th percentiles of the dollar volume for common NYSE-listed stocks. Group 1 (Group 10) contains stocks with the lowest (highest) trading volume. The standard errors are clustered at weekly levels for 17 industries. The sample ranges from January 2001 to December 2005.

Figure 3: Price Impact and Spread across 10 Volume Groups

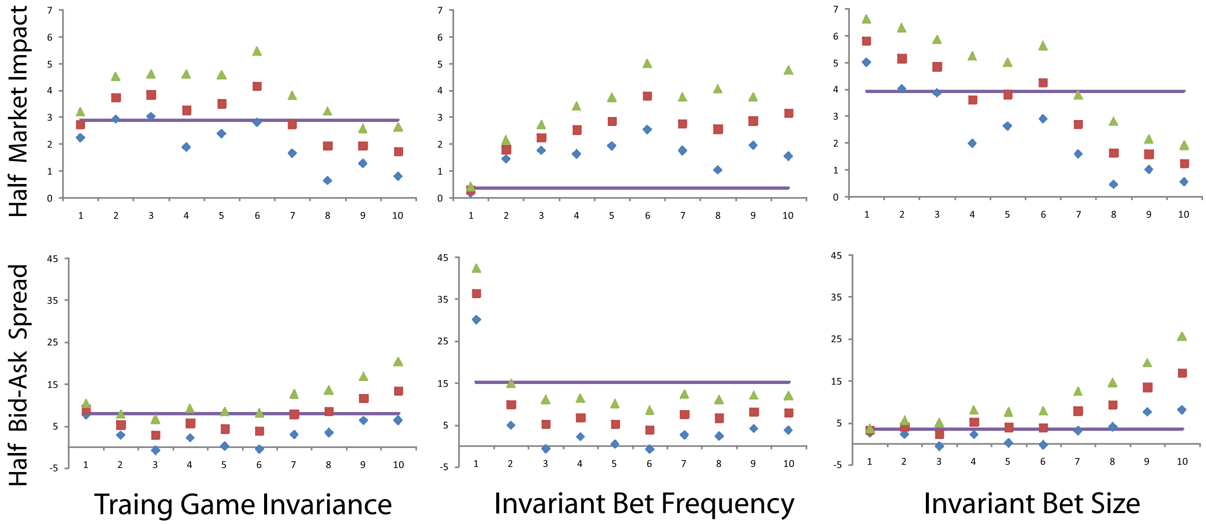

Figure graphs the estimates of half price impact $\frac{1}{2}\bar{\lambda}_j$ (top plots) and half effective spread $\frac{1}{2}\bar{k}_j$ (bottom plots) with the 95%-confidence intervals for 10 volume groups and for three proposed models from the regression $Y_i = \left(\sum_{j=1}^{10} \mathbb{I}_{j,i} \times \frac{1}{2}\lambda_j\right) \times \left[\frac{W_i}{W_*}\right]^{\alpha_0} \frac{X_i}{(0.01)V_i} + \left(\sum_{j=1}^{10} \mathbb{I}_{j,i} \times \frac{1}{2}k_j\right) \times \frac{(X_{omt,i}+X_{ec,i})}{X_i} \left[\frac{W_i}{W_*}\right]^{\alpha_1} + \tilde{\epsilon}$ as well as unconditional estimates from Table 7. In Model of Trading Game Invariance, $\alpha_0 = 1/3, \alpha_1 = -1/3$. In Model of Invariant Bet Frequency, $\alpha_0 = 0, \alpha_1 = 0$. In Model of Invariant Bet Size, $\alpha_0 = 1/2, \alpha_1 = -1/2$. Each observation corresponds to order i. $\mathbb{I}_{j,i}$ is an indicator equal to one if order i is executed in a stock from group j. $\bar{\lambda}_j/2$ estimates in basis points the market impact costs of a trade of one percent of average daily volume in a benchmark stock for volume group j, and $\bar{k}_j/2$ estimates in basis points the effective spread cost for volume group j. Y_i is the implementation shortfall in basis points calculated as $\frac{\mathbb{I}_{BS,i}(P_{ex,i}-P_{0,i})}{P_{0,i}} 10^4 \frac{(0.02)}{\sigma_{r,i}}$, where $\mathbb{I}_{BS,i}$ is a buy/sell indicator, $P_{ex,i}$ is the average execution price, $P_{0,i}$ is the pre-transition price, $\sigma_{r,i}$ is the expected daily volatility estimated as $\sigma^e_{i,t-1}$. The term $(0.02)/\sigma_{r,i}$ adjusts for heteroscedasticity. The trading activity W_i is the product of expected volatility $\sigma_{r,i}$, benchmark price $P_{0,i}$, and expected volume V_i measured as last month's average daily volume. The scaling constant $W_* = (0.02)(40)(10^6)$ corresponds to W_i for the benchmark stock with volatility of 0.02, price \$40 per share, and trading volume of one million shares per day. X_i is the number of shares in the order with $X_{omt,i}$ executed in open market and $X_{ec,i}$ executed in external crossing networks. The term $X_i/(0.01)V_i$ is the size of the trade relative to expected daily volume. Volume groups are based on the pre-transition dollar trading volume with thresholds corresponding to 30th, 50th, 60th, 70th, 75th, 80th, 85th, 90th, and 95th percentiles of the dollar volume for common NYSE-listed stocks. Group 1 (Group 10) contains stocks with the lowest (highest) trading volume. The standard errors are clustered at weekly levels for 17 industries. The sample ranges from January 2001 to December 2005.