Short-Term Trading Skill: An Analysis of Investor Heterogeneity and Execution Quality*

Mehmet Sağlam†
Lindner College of Business
University of Cincinnati
email: mehmet.saglam@uc.edu

Ciamac C. Moallemi
Graduate School of Business
Columbia University
email: ciamac@gsb.columbia.edu

Michael G. Sotiropoulos
Princeton University
email: msotirop@princeton.edu

Draft: January 2016

Abstract

We examine short-horizon return predictability using a novel proprietary dataset of institutional traders with known identities. We estimate investor-specific short-term trading skill and find that there is pronounced heterogeneity in predicting short-term returns among institutional investors. Incorporating short-term predictive ability, our model explains much higher fraction of variation in asset returns. Ignoring the heterogeneity in short-term trading skill has major implications in quantifying execution quality. A simple trading strategy exploiting our skill estimates yields economically and statistically significant abnormal return. Finally, we provide an *ex ante* metric to identify skilled and unskilled investors that can be used to estimate execution costs.

Keywords: Short-term Return Predictability, Institutional Trading, Execution Quality, Algorithmic Trading

JEL Classification: G11, G14, G24.

---

*We are grateful for helpful comments from Robert Battalio, Alex Borisov, Zhi Da, Larry Glosten, Hui Guo, Brian Hatch, Gerry Tsoukalas, Kumar Venkataraman, Haoxiang Zhu and seminar participants at Columbia University, 2012 INFORMS conference, Princeton University and University of Cincinnati. Sağlam acknowledges support from Deming Doctoral Fellowship. Moallemi acknowledges support from NSF Grant CMMI-1235023.

†Corresponding author: Mehmet Sağlam, University of Cincinnati, 2925 Campus Green Drive, Cincinnati, OH, 45221-0195, Phone: (513) 556-9108; Fax: (513) 556-0979; mail: mehmet.saglam@uc.edu

1. Introduction

There is ample evidence that excess stock returns are predictable at various horizons by macroeconomic and firm-level characteristics such as dividend-price and book-to-market ratios, short-term rates, aggregate volatility, and lagged returns. Indeed, there is ongoing active research in uncovering predictive variables or proposing trading strategies generating abnormal returns against standard asset-pricing models. This literature is motivated by the fact that investors can exploit predictable returns when making portfolio decisions dynamically. For example, Johannes et al. (2014) find statistically and economically significant benefits for investors using models of return predictability.

In addition to the documented sources of return predictability, some investors may have private information about the fundamental value of the asset. It is unlikely that this informational advantage is due to having access to non-public material information as a corporate insider but, rather, some investors may just be more skilled in processing short-term information flows to identify under- or over-valued stocks. For example, using a large set of institutional trading data, Yan and Zhang (2009) show that trades of short-term institutional investors are positively correlated with future stock returns. Similarly, Diether et al. (2009) examine daily short-sale trading activity and argue that short-sellers can detect when the asset prices deviate from their fundamental value.

These examples show that, exploiting particular return-predictive signals, some investors will be able to forecast short-term price movements. The literature on the evidence of informed trading by institutional investors mostly focuses on these types of skilled investors. However, not all investors are motivated by short-term goals. Investors will also undertake certain trading strategies that are idiosyncratically dependent on their own investment objectives, style, and horizon, and may end up employing trading strategies that are at odds with short-term return predicting signals. As an extreme case, consider a fund manager with a value-style investment view and a long-horizon performance benchmark. Asness et al. (2013) document that value and momentum signals are negatively correlated and have different time horizons, hence a value investor may systematically trade in a way that is opposite to short-term predictable returns generated through momentum effects.

In this paper, we analyze a large data set of intraday institutional trading data with masked investor identities, and decompose the performance of trades into two components: (i) an investor-specific "trading skill" component which captures the timing of the decision to trade relative to favorable or unfavorable short-term price movements; and (ii) a market impact or price impact component, which measures investor-independent execution costs. We document the presence of *skilled* short-term investors in our data set. They are identified as the ones that systematically buy (resp., sell) an asset during a period when the asset return is positive (resp., negative). In addition, we find strong evidence for the presence of *unskilled* investors in our data. These systematically decide to buy (resp., sell) the asset during a trading interval when the asset return is negative (resp., positive) on average. This heterogeneity in short-term trading skill has important implications for measures of execution costs, such as the implementation shortfall, introduced by Perold (1988). Since skilled traders correctly predict short-run future returns, the cost of their trades appear high when compared to the trades of a benchmark noise trader. Similarly, because unskilled traders make trading decisions that are systematically opposite to short-term returns, the execution cost of their trades appear low when compared to a noise trader. As a result, measured execution costs may not be an unbiased estimate of the true cost of trading, which has been the crucial measure of market quality assessment in the literature. For example, a number of earlier studies (e.g., Huang and Stoll (1996), Bessembinder and Kaufman (1997a), Bessembinder and Kaufman (1997b)) utilize execution costs to compare execution quality differences between NYSE and NASDAQ. Similarly, in order to improve the transparency on market quality, the Securities and Exchange Commission (SEC) adopted Rule 605 on November 15, 2000, which requires market centers to make monthly public disclosure of certain execution costs. Given this regulatory emphasis, execution costs have also been a popular comparison metric with various changes in market structure. O'Hara and Ye (2011) use execution costs to test the impact of market fragmentation on market quality. Brogaard et al. (2014) and Tong (2015) examine the impact of high-frequency trading on the executions costs of institutional investors. All of these studies motivate the significance of obtaining accurate measures of execution quality for the use of brokers and policy-makers.

There is also evidence from both theoretical and empirical literature that heterogeneity in trading skill may affect the choice of selecting different venues for trading needs. Zhu (2014) and Iyer et al. (2015) argue that informed traders strategically choose the lit markets for their execution needs, whereas dark pools are relatively more attractive to uninformed traders. Consequently, a naive execution cost analysis that does not take this into account may systematically suggest that (assuming all else is equal) dark pools have better execution quality. Similarly, there is strong evidence that short-term information may affect the choice of limit orders versus market orders, or the choice of a high-rebate or low-rebate trading venue (e.g., Kaniel and Liu (2006), Maglaras et al. (2012), Collin-Dufresne and Fos (2015)). Thus, an execution cost analysis across trading venues controlling for all other effects but not investor heterogeneity should similarly be systematically biased.

In this paper, we are interested in jointly estimating an investor-dependent short-term trading skill and an investor-independent measure of execution costs. To our knowledge, the effects of heterogeneous short-term predictive ability have been largely ignored in prior execution cost studies. In general, these effects are much more difficult to model *ex ante*, since short-term predictions cannot be observed directly. In a typical algorithmic trading situation, where an investor executes a large order through an algorithm provided on an agency basis by a broker, the investor rarely communicates their short-term price views directly to the executing broker. Instead, investors might implicitly express their alpha view by choosing the asset, direction, and time of the trade, and by adjusting the parameters of the broker's trading algorithm. By not accounting for short-term trading skill, any subsequent transaction cost analysis may misestimate the price impact associated with the investor's trades.

For this purpose, we propose a model to attribute the asset returns observed during the execution of a large order between the short-term predictive skill of the investor and the price impact of the resulting trades. Specifically, we consider short-term trading skill as a characteristic of the investor. Besides the usual price impact factors such as the relative size of the order, speed and volatility, the model introduces the investor's short-term predictive ability in the form of risk-adjusted performance metric as in the typical usage with Sharpe ratio. We do not impose any *a priori* grouping of investors into categories that might explain their predictive ability, such as institutional investors, quantitative funds, or retail investors. The risk-adjusted measure of short-term trading performance allows our model to capture the dependence of future price movements on the mere desire of an investor to trade a specific asset at a specific point in time.

We estimate our proposed model on a unique and proprietary historical data set consisting of a large sample of intraday equity execution data along with masked investor identifiers, obtained from a large broker who provides algorithmic trading services. We analyze our estimation results for robustness, and our contributions can be summarized as follows:

1. There is strong evidence for investor heterogeneity in short-term trading skill. We find that approximately one third of the investors are systematically skilled or unskilled relative to the rest. In other words, ability to predict short-term price changes may be a significant motivation for many investors in our sample to trade a specific asset at a specific point in time. A permutation test provides further evidence that the numbers of skilled and unskilled investors are abnormally high.

2. We construct a simple out-of-sample trading strategy based on skill estimates and find that this trading strategy generates significant abnormal returns of 18.8% at the annualized level when benchmarked against the Carhart (1997) four-factor model. In other words, our classification of skilled and unskilled traders is consistent with their ability to predict future returns in the short-term.

3. Short-term trading skill significantly increases the power of the model in explaining the variation of returns relative to arrival price. In fact, including investor specific skill variables improves the R^2 of the model relative to a model that only considers the price impact of orders by an order of magnitude, from 0.5% to 10%. In other words, the identity of an investor who wishes to trade is highly predictive of future price movements relative to considering only the orders the investor places. Moreover, ignoring investor identity results in systematic misestimation of the price impact of trades. Our results are robust to alternative model specifications and can actually predict out-of-sample returns with skill estimates.

4. Short-term skill has statistically significant dependence on asset- or order-specific characteristics. Short-term trading skill is positively correlated with *ex ante* short-term momentum and order size relative to normal trading volume. Buy (resp., sell) orders placed by investors with higher skill tend to occur when the asset has higher (resp., lower) short-term cumulative past returns and the order size is relatively larger (resp., smaller) when expressed as a percentage

of average daily volume. Moreover, the orders of skilled traders tend to be associated with shorter execution duration, higher participation rate and larger bid-ask spreads. These are all hallmarks of more urgent and more informed trading. Explanation of skill with short-term momentum is also consistent with the recent evidence of Moskowitz et al. (2012) on return predictability with time-series momentum.

5. We find that out-of-sample execution costs have major statistical dependence on our skill estimates. Expected execution cost difference between short-term skilled and unskilled traders is 25 bps which is economically substantial. We find that *ex ante* predictive variables of short-term trading skill – short-term momentum and relative size – can explain more than 50% of this cost differential.

In current literature, there has been little evidence of the cross-sectional structure of short-term trading skill across a universe of institutional investors. Our paper proposes a methodology to identify such investor behavior. We demonstrate that short-term predictive ability is very heterogeneous among an institutional investor base. This is consistent with theoretical agent-based microstructure models where information asymmetry provides a major motivation to trade. Moreover, while the literature focuses on informed and uninformed investors, our results reveal the presence of another type that systematically places orders in the opposite direction of short-term future returns.

From a policymaker perspective, our results illustrate that mere comparison of execution costs cannot be a standalone measure of execution quality. Venues populated either with skilled or unskilled traders may have misestimated measures of execution quality if the heterogeneity in short-term trading skill is ignored. Consequently, these biased estimates may not lead to an optimal policy recommendation.

From a practical perspective, moreover, our results illustrate that incorporating short-term trading skill is important in the estimation of execution costs and, in particular, of price impact. Ignoring skill heterogeneity results in models that have both much lower predictive ability and systematically biased estimates of price impact, which may often be conflated with skill. Eliminating this bias may result in improved decision-making throughout the trading process. In the pre-trade phase, for example, more accurate transaction cost estimates will result in better portfolio construction. During trade execution, accounting for short-term trading skill will allow brokers to tailor their trading algorithms on an investor-by-investor basis and achieve better execution results. Finally, our predictive variables for short-term trading skill can be utilized in the absence of investor identities or limited execution data availability so that execution costs will be estimated more accurately in the pre-trade phase.

The rest of the paper is organized as follows: in Section 1.1, we present a brief literature review. In Section 2, we set up the underlying statistical model. Section 3 describes our experimental study, while Section 4 contains our model estimation and analysis. Sections 5 and 6 examine the robustness of our results and provide strong evidence for our interpretations with investor heterogeneity in short-term trading skill. Section 7 presents the dependence of investor skill on *ex ante* and *ex post* asset- and order-specific characteristics. Finally, Section 8 discusses the implications of investor heterogeneity on measures of execution quality. We conclude in Section 9.

1.1. Literature Review

Our paper is related to two main strands of the literature: studying skill in institutional trading and estimating the price impact of trading activity.

A large literature on institutional trading activity addresses the question whether institutional investors are informed. Gompers and Metrick (2001) find that there is positive relationship between institutional ownership and future stock returns. Yan and Zhang (2009) argue that this relationship is driven by short-horizon institutions. Using a more high-frequency data, Puckett and Yan (2011) find that institutional investors are skilled even after accounting for trading costs. In a more recent study using news analytics, Hendershott et al. (2015) find that institutional investors are informed and their trading direction can predict the sentiment of the future news. There are also a number of studies that document skill in the general context of fund management (e.g., Cohen et al. (2005), Kacperczyk et al. (2005), Mamaysky et al. (2008)). On the other hand, Anand et al. (2012) document that institutional trading desks have persistent trading costs – institutions that have low trading costs continue to have low trading costs over time. Our paper is related to this literature but focuses on studying the heterogeneity in short-term trading skill. In terms of documenting unskilled short-term investors, our results also resemble the underperformance of fund managers after accounting for management fees as in Wermers (2000).

The relationship between trading activity and asset prices in financial markets has been an important question in the economic microstructure literature for several decades. The theoretical origins of price impact arise from the presence of informed traders as, for example, in the celebrated models of Kyle (1985) or Glosten and Milgrom (1985). As a result, a line of literature has emerged focusing on the empirical analysis of the impact of trades on prices, motivated by the economic question of understanding the role of information asymmetry in markets. This work is nicely summarized by Hasbrouck (2007) and it is still actively pursued, see Easley et al. (2012).

More recently, however, with the rise of electronic and algorithmic trading, a new line of literature has emerged. Motivated by the concerns of practitioners, this literature focuses on the decision problem faced by an investor seeking to algorithmically spread his trades out over time, in order to minimize execution costs. A key ingredient in such algorithmic trading is the estimation of the effect of a sequence of "child" orders placed by an algorithm executing an investor's "parent" order on the asset price across future time horizons. The most notable early works here are those of Bertsimas and Lo (1998) and Almgren et al. (2005). More broadly, Bouchaud et al. (2008) provides a summary of theoretical and empirical results on models which predict the impact of trades on prices, bid-ask spreads, and other market dynamics over time. They theorize that much of these dynamics can be explained by the presence of algorithmic traders strategically spreading their orders across time.

A closely related question is the estimation of overall transaction costs for large block trades (e.g., Keim and Madhavan (1996), Almgren (2008)). These cost functions play an important role in portfolio optimization and other pre-trade analysis. Obizhaeva (2009) estimates such trading cost functions using a data set of large portfolio transitions. Kyle and Obizhaeva (2014) provide a theoretical model that seeks to explain the cross-sectional variation of trading costs across a universe of stocks. Our work extends this line of inquiry by explicitly including investor identity as a predictive factor of order execution costs.

Finally, our paper is related to the analysis of transaction costs with respect to different investment strategies. Using equity executions from 21 institutional traders, Keim and Madhavan (1997) find that total trading costs for a technical-style investment strategy is higher compared with value-style investment strategy. Intuitively, they relate this finding to the differences in aggressiveness as value investors trade patiently via worked orders. Similarly, using transactions of 37 money managers, Chan and Lakonishok (1995) find that growth-oriented strategies incur higher transaction costs due to differences in their demands for immediacy. Controlling for differences in trading schedules or demands for immediacy, our paper complements these studies focusing on a fundamentally different theme. In our model, heterogeneity across investors does not stem from better trading schedules as this is ultimately controlled by the algorithm of the broker, but from differences in skill levels for predicting short-term returns. We also further control for differences in demand for immediacy via the order's participation rate and duration.

2. The Model

We consider a population of J investors sending a total of N orders to an executing broker. The mapping $i \xrightarrow{c} j$ identifies order i, $i = 1 \ldots N$ as belonging to investor $j = c(i)$, with $j = 1 \ldots J$. Each order is for a quantity of Q_i shares of an asset, with $Q_i > 0$ ($Q_i < 0$) for buy (sell) orders, respectively. Each order also has an execution duration of T_i, measured as a fraction of the trading day. We define the participation rate $\rho_i \triangleq |Q_i|/V_i$, where V_i is the total market volume traded within the interval T_i. The arrival price $P_{i,0}$ is the last traded price prior to the order's arrival and the terminal price is the last execution price P_{i,T_i}. In our model, given order i, we consider the expected return of the asset over the execution interval T_i, that is, $\log(P_{i,T_i}/P_{i,0})$. We posit that this return is driven by two predictable effects. The first effect is *short-term trading skill*. Some investors will be able to predict short-term asset returns using models of return predictability, such as time-series momentum as documented by Moskowitz et al. (2012). Conditional on the arrival of a buy (resp., sell) order i, we expect a return in the asset price of $\alpha_{c(i)}\sigma_i\sqrt{T_i}$ over the execution interval T_i. Here, The coefficient $\alpha_{c(i)}$ represents the short-term predictive ability of investor $c(i)$ and σ_i is the daily volatility of the mid-quote of the asset price, typically estimated as an average of daily volatilities over the prior month. $\alpha_{c(i)}$ can be positive, zero, or negative which can be interpreted as skilled, unidentified or unskilled respectively. Note that this predictive ability is parameterized in a risk-adjusted fashion, i.e. we assume that each particular investor has a constant short-term Sharpe ratio or information ratio over all of the trades.[1]

The second effect is *price impact*, or, the direct effect of the trades placed on behalf of the investor. The price impact for order i is given by $\lambda\sigma_i\sqrt{T_i}h(\rho_i)$. Here, λ is a (broker-specific) price impact coefficient. We normalize the price impact component with the volatility of the asset during execution horizon captured by $\sigma_i\sqrt{T_i}$ so that we represent the impact as a fraction of the typical movement of the stock return. Transaction cost models sometimes include a bid-offer spread term to incorporate stock-specific liquidity costs. However, since we are modeling *returns* over the execution horizon as opposed to the average cost, we did not include spreads in our baseline specification. In Section 5.3, we include an additional spread component in the price impact specification and our findings remain largely unchanged. Our price-impact assumption is consistent with the literature (e.g., Almgren et al. (2005)). On the theoretical front, Keim and Madhavan (1996) derive that price impact is a concave function of the trade size. Similarly, Chacko et al. (2008) also finds that expected price impact is proportional to the volatility and empirically validates this claim.

The price impact function $h(\cdot)$ captures the effect of the participation rate (or trading speed) on price. The idea here is that orders executed with a higher participation rate will have a larger price impact. In order to illustrate the robustness of our results with respect to our price impact formulation, we will consider two explicit forms for the price impact function: a linear price impact function, i.e.,

$$h(\rho_i) \triangleq \rho_i,$$

or a square root price impact function, i.e.,

$$h(\rho_i) \triangleq \rho_i^{1/2}.$$

The choice of sublinear price impact has been extensively studied both theoretically and empirically. There is a long line of literature supporting the choice for a square root price impact law.

[1]Here, we ignore the role of a benchmark risk-free return in the definition of a Sharpe ratio, i.e. we do not consider excess returns. This is reasonable since the risk-free return is effectively zero over the intraday time horizons of interest.

Further, note that we do not scale Sharpe ratio with the square root of the investment horizon (as typically done with longer horizons in asset management). In Section 5.2, we explore an alternative specification where we define $\alpha_{c(i)}$ to be a daily Sharpe ratio that is scaled by the square root of the length of the execution horizon, and find that both models provide very similar findings.

For example, Chacko et al. (2008) provides empirical evidence that the expected price impact is proportional to the square root of the quantity traded. Using a large sample of US equity trades, Almgren et al. (2005) also estimate the exponent to be very close to 0.5. This exponent is also consistent with the well-known Barra model for market impact costs outlined in Torre and Ferrari (1998).

Putting everything together, we assume that the sign-adjusted log-return of an order relative to the arrival price and over the execution horizon can be expressed as an additive model of the form

$$\operatorname{sgn}(Q_i)\log\left(\frac{P_{i,T_i}}{P_{i,0}}\right) = \alpha_{c(i)}\sigma_i\sqrt{T_i} + \lambda\sigma_i\sqrt{T_i}h(\rho_i) + \epsilon_i\,, \tag{1}$$

with ϵ_i having a mean of zero and variance of ν_i^2.

[Insert Figure 1 here]

Figure 1 illustrates the components of short-term trading skill and price impact in our model when an investor sells her shares in a stock. Here, alpha term refers to $\alpha_{c(i)}\sigma_i\sqrt{T_i}$ which can also be negative as previously discussed.

Explicit forms of the impact function $h(\cdot)$ fully specify the model as a linear regression of the risk-normalized interval return against short-term trading skill and broker impact. For example, with a linear price impact function we obtain

$$\operatorname{sgn}(Q_i)\log\left(\frac{P_{i,T_i}}{P_{i,0}}\right) = \beta_0 + \sigma_i\sqrt{T_i}\sum_{j=1}^{J}\mathbb{I}_{\{c(i)=j\}}\alpha_{c(i)} + \lambda\sigma_i\sqrt{T_i}\rho_i + \epsilon_i\,, \tag{2}$$

with $\mathbb{I}$ the indicator function. Likewise, a square-root impact function leads to the model

$$\operatorname{sgn}(Q_i)\log\left(\frac{P_{i,T_i}}{P_{i,0}}\right) = \beta_0 + \sigma_i\sqrt{T_i}\sum_{j=1}^{J}\mathbb{I}_{\{c(i)=j\}}\alpha_{c(i)} + \lambda\sigma_i\sqrt{T_i}\rho_i^{1/2} + \epsilon_i\,. \tag{3}$$

When fitted on a historical sample of short-term execution returns, the above models (2) and (3) identify the short-term trading of each investor j, along with the impact coefficient.[2]. We proceed

[2]The above models could also be expressed in terms of the arrival slippage $\log\left(\bar{P}_i/P_{i,0}\right)$ instead of the interval

by presenting our data set and model estimation results in Section 4.

3. Data

For our empirical study we use a novel proprietary execution data from the historical order databases of a large investment bank ("The Bank"). The orders originate from a diverse pool of investors, such as institutional portfolio managers, quantitative investment funds, internal trading desks and retail customers. Our data set consists of two widely used algorithmic execution strategies, the volume weighted average price (VWAP) and the percentage of volume (POV). The VWAP algorithm aims to achieve an average execution price that is as close as possible to the volume weighted average price over the execution horizon. The main objective of the POV algorithm is to have constant participation rate in the market within the execution interval. VWAP and POV have relatively small discretion on opportunistically speeding up or slowing down the execution so the aggressiveness of the execution is mainly controlled by the investor choosing a particular urgency level in the pre-trade phase. With VWAP and POV algorithms, we eliminate any potential broker-specific effects such as the usage of trading signals and market events that drive more opportunistic algorithms. Furthermore, we also avoid any biases that can occur with endogenous selection of the algorithm itself. For example, if we were to have more sophisticated algorithms in our data set, one might argue that skilled traders may actually be just better in choosing algorithms. However, with VWAP or POV, an investor can only have superior short-term trading skill by starting the execution at a particular time and communicating a desired execution horizon to the broker via adjusting the urgency parameters of the algorithm.

This proprietary data set provides a rich set of attributes. For each order i we have access to the following: investor identity tag,[3] $c(i)$, ticker of the traded stock, order size, Q_i, order side (buy/sell), $\operatorname{sgn}(Q_i)$, execution duration, T_i, participation rate, ρ_i, average volatility of the stock over the last 20 trading days, σ_i, the percentage return over the execution interval, $P_{i,T_i}/P_{i,0} - 1$.

return $\log(P_{i,T_i}/P_{i,0})$, where $\bar{P}_i$ is the average execution price of the i-th order. This would lead to an approximate rescaling of the coefficients $\alpha_{c(i)}$ and λ by a factor of $1/2$. The execution algorithms considered here trade at constant participation rate. Therefore, the execution price $\bar{P}_i$ is close to the realized interval VWAP, and for a price path with a constant drift, the VWAP return is half of the interval return.

[3] Investors are identified by numerical aliases to protect anonymity.

These data allows us to fully estimate the model of Section 2. In addition, our data include the daily, average (i.e., over the last 20 trading days[4]) and interval (i.e., during the execution horizon) proportional bid-ask spread, mid-quote volatility and traded volume for each stock.

We use a restricted subset of the execution data, defined by the following selection criteria:

- The trading period is from January 2011 to June 2012, inclusive.
- The asset universe consists of the S&P 500 stocks. We focus on highly liquid stocks to focus on the differences on short-term predictive ability as a result of following certain set of strategies. For this set of stocks, it is hard to have an investor trading on an insider information.
- Orders come from active investors only: an investor is considered active if he has at least 100 and at most 500 orders within the period of study. This is to prevent any specific investor from having major influence on our results.
- All orders have been fully filled without intermediate replacements or cancellations.
- The execution duration is greater than 5 minutes but no longer than a full trading day, with no participation in opening/closing auctions. This is to avoid any short term effects from market orders and auctions.

Using the above criteria, our final sample consists of 63,379 executions coming from 30,438 buy and 32,941 sell orders. The trading algorithms used are 41,339 VWAP and 22,040 POV. The orders came from a set of 293 active investors, with 216.3 orders per investor and 168 orders per trading day on average. The highest number of executions on a single stock is 454 which corresponds to 0.71% of all executions. Table 1 provides additional summary statistics for our data sample.

[Insert Table 1 here]

The average percentage return realized during the execution interval is 0.3 bps. We observe that bid-ask spread and volatility lay in a tight range. More than half of the executions have a bid-ask spread between 2 bps and 5 bps and a mid-quote annualized volatility between 15% and

[4]Throughout the analysis we will refer to this measure as the average over the past-month for conciseness as there are 21 trading days in each month on average.

27%. The mean duration of the executions is a little less than 2.5 hours. Finally, we have a wide range of participation rates across executions with an average (median) of 6.44% (1.59%).

4. Model Estimation Results

We analyze the execution data using our full model with two price impact specifications, linear price impact as in (2) and square root price impact as in (3). We also fit a reduced model to the data by ignoring trading skill terms, i.e., $\alpha_j = 0$.

$$\operatorname{sgn}(Q_i) \log \left(\frac{P_{i,T_i}}{P_{i,0}} \right) = \beta_0 + \lambda^{\mathsf{base}} \sigma_i \sqrt{T_i} \rho_i^{\gamma} + \epsilon_i \, , \tag{4}$$

where $\gamma = 1$ for the linear model and $\gamma = \frac{1}{2}$ for the square-root model. We use the superscripts, base, to emphasize the difference between the reduced and full models.

We are concerned with heteroscedasticity, contemporaneous correlation across stocks, and autocorrelation within each stock and adjust our standard errors by clustering on calendar day and stock throughout the analysis as suggested by Petersen (2009). The regression results are summarized in Table 2.

[Insert Table 2 here]

Table 2 suggests a number of interesting observations. First, consider the price impact parameter, λ. In all cases, the estimate of λ is statistically significant. The square model in the absence of trading skill corresponds to the well-known Barra market impact model as outlined in Torre and Ferrari (1998). Here, our estimate of λ is of order unity — this is consistent with the prior literature. Yet, we observe that without accounting for investor's short-term trading skill level, price impact by itself has a very low explanatory power with the maximum adjusted R^2 of 0.52% from both models. The linear price impact model has relatively better fit than the square root price impact model, even though the difference is very small. However, the inclusion of the short-term trading skill term substantially increases the goodness of fit, leading to an adjusted R^2 of approximately 10%. This significant difference illustrates that the variation in short-term returns can be explained much better when the systematic short-term trading skill of the investor is acknowledged.

[Insert Table 3 here]

Moreover, if we ignore the predictive abilities of the investors, we observe that price impact is misestimated. If the price impact is linear in participation rate, then accounting for alpha view of the investors reduces the price impact coefficient by approximately 20%. This is also observed for the square-root model, but to a lesser degree[5]. The standard errors do not allow us to conclude that the difference between $\hat{\lambda}$ and $\hat{\lambda}^{\mathsf{base}}$ is statistically different. For this reason, we use a bootstrapping analysis in which we construct $1,000$ datasets each with $10,000$ randomly assigned executions and estimate the full and the reduced models. We find that the difference between $\hat{\lambda}$ and $\hat{\lambda}^{\mathsf{base}}$ is highly statistically significant in both price-impact specifications. Table 3 illustrates these findings formally.

We observe that there is significant investor heterogeneity in predicting short-term returns. We label an investor as *skilled* (resp, *unskilled*) if his short-term trading skill estimate is positive (resp., negative) and is statistically significant under 10% level. We label the remaining investors as *unidentified*. Table 2 reports that with the linear price impact model, 49 out of 293 investors are *skilled*, while 48 investors are *unskilled* and 196 investors are *unidentified*. With the square root price impact model, the alpha estimates slightly drop as the square-root model puts more weight on the price impact unilaterally. Figure 4 illustrates this drop in Sharpe ratio estimates graphically. In this case, we obtain that 36 investors are *skilled* whereas 63 are *unskilled*. In both models, approximately one third of our investor universe are *skilled* or *unskilled*.

[Insert Table 4 here]

These findings do not necessarily imply that the number of *skilled* or *unskilled* investors in our sample is abnormal. However, we can implement a permutation test to assess the statistical significance of these numbers. Our dataset allows us to generate the empirical distribution of the number of *skilled* and *unskilled* traders under the null hypothesis that investor identifiers are unrelated to log-returns during the execution horizon. We create $10,000$ different samples of our execution dataset by permuting the investor identifiers across executions. Each investor has still

[5] This result suggests that in our sample short-term skilled trader activity is relatively higher when compared with unskilled trader activity.

the same number of assigned executions but in this case all executions are reshuffled randomly among all the investors. We then estimate the regression coefficients from each resampled dataset and compute the empirical distribution for the desired parameters: the numbers of *skilled* and *unskilled* investors along with the price impact coefficient, $\hat{\lambda}$. Table 4 illustrates the p-values of our original estimates of the numbers of *skilled* and *unskilled* investors and the price impact coefficient. In both linear and square-root models, we find strong evidence that the estimated number of *skilled* or *unskilled* investors is indeed abnormally high. Similarly, for both models, the estimated price impact coefficient, $\hat{\lambda}$, is lower than the minimum coefficient computed in all of the $10,000$ resampled datasets highlighting the divergence from the null hypothesis.

[Insert Figure 2, Figure 3, and Figure 4 here]

We now discuss the magnitudes of the estimated short-term predictive abilities. Figure 2 shows histograms of investor skill estimates both for the linear and square root price impact model specifications, including statistically insignificant estimates. Figure 3 shows the histograms of the alpha estimates which are statistically significant. We observe that estimates which are small in absolute value are likely to be insignificant. We report the investor short-term trading skill estimates as annualized Sharpe ratios. In the linear model, the range of the estimated Sharpe ratios is between -27.7 and 14.3 with the sample mean (median) of -0.59 (-0.40). We find that the distribution of skill estimates have negative skew and large kurtosis suggesting that the distribution of the skill estimates are asymmetric and non-normal. The Sharpe ratio estimates arising from our models are much larger than typical Sharpe ratios observed in the traditional asset management industry. However, current empirical literature reports similar Sharpe ratio estimates over short, intraday investment horizons (e.g., for high-frequency traders). For example, Clark-Joseph (2013) estimates that annualized Sharpe ratios of high-frequency traders are in the neighborhood of 10 to 11. Baron et al. (2013) report that the average high-frequency trader Sharpe ratio in their data set is 9.2.

In summary, we observe that accounting for investor heterogeneity in short-term trading skill explains much higher variation of asset returns during an execution. In the absence of a trading skill term, we observe that the price impact of a trade is highly overestimated. This suggests that the usual practice of ignoring investor specific view introduces a significant bias in price impact estimates. We further investigate this implication in the context of execution costs in Section 8, and find that the standard measure of execution cost, implementation shortfall, depends on short-term predictive ability.

Our findings suggest that, at the individual investor level, there is substantial variation in short-term predictive ability. Roughly, one third of the investors have statistically significant Sharpe ratio estimates. On the other hand, we observe that half of these investors make systematically wrong bets in the short-term. Note that these statistics do not imply that investors in this group are losing money in the long-run. These investors are possibly long-term investors who are not exploiting the short-term predictability of asset returns. In Section 7, we find further evidence that skilled investors utilize short-term momentum in asset returns.

5. Robustness Tests

In this section, we assess the robustness of our results in five ways. First, in order to control for the possibility of over-fitting, we assess the robustness of our model predictions on out-of-sample data. Using our estimated model in-sample, we predict out-of-sample short-term returns over an execution. Second, we explore a different alpha specification that scales with the square root of the execution horizon. Third, we consider an alternative price impact specification with an additional spread component. Fourth, we test the robustness of our findings when execution horizon is largely the same across orders and is not in the subset of slowest or fastest executions. Fifth, we consider an alternative specification that incorporates the market return over the execution time horizon.

5.1. Out-of-Sample Predictions

Our model specifications in equations (2) and (3) contain a number of parameters, namely, one for each investor. In order to eliminate the possibility of over-fitting, in this section, we consider a cross-validation experiment that illustrates the ability of our model to predict out-of-sample execution returns.

First, we divide the data into two parts: in-sample data and out-of-sample data. We perform this by randomly allocating half of each investor's executions into the in-sample data set and the remaining ones into the out-of-sample data set. We then estimate the model parameters by running the regressions specified in equations (2) and (3) using only the in-sample data.

[Insert Table 5 and Figure 5 here]

Table 5 illustrates the regression results for the in-sample data set. The estimated regression coefficients for price impact are very similar to those obtained using all the data. For example, using the linear price impact specification, the price impact estimate, $\hat{\lambda}$, is 1.79 whereas using the complete data, the estimate is 1.81. Similarly, we observe that investor skill estimates are also very stable. Figure 5 compares skill estimates between the in-sample and the complete data sets. In both price impact models, these are very close to each other, implying the robustness of the estimates.

[Insert Table 6 here]

Using the skill and price impact estimates obtained from the in-sample data, we can test whether our model can explain out-of-sample execution returns. Table 6 illustrates the root mean squared prediction error (rMSPE) and R^2 estimates both in-sample and out-of-sample. We observe that in-sample and out-of-sample mean-squared errors are very close to each other. Similarly, we obtain an out-of-sample R^2 of more than 8.1% in both price impact models suggesting that our regression model does not suffer from over-fitting. Both of these findings emphasize that our model has out-of-sample predictive power. In Section 8, we will again utilize these in-sample and out-of-sample data to study the implications for execution quality.

5.2. Robustness in Alpha Specification

Sharpe ratio is typically defined over a reference time horizon, and is often scaled with the square root of the investment horizon when it is projected across different horizons. In our model, however, we assumed that the Sharpe ratio is held constant, independent of the execution horizon. We can also consider the alternative model. In this specification, investor j has an expected return of $\sigma_i\sqrt{T_i}\alpha_j\sqrt{T_i}$ when he trades ith stock during the trading horizon, T_i. Consequently, our skill estimation models are given by

$$
\operatorname{sgn}(Q_i)\log\left(\frac{P_{i,T_i}}{P_{i,0}}\right) = \sigma_i T_i \sum_{j=1}^{J} \mathbb{I}_{\{c(i)=j\}}\alpha_{c(i)} + \lambda\sigma_i\sqrt{T_i}\rho_i^{\gamma} + \epsilon_i\,, \tag{5}
$$

where $\gamma = 1$ for the linear model and $\gamma = \frac{1}{2}$ for the square-root model.

[Insert Table 7 here]

Table 7 illustrates the regression results for the model presented in Equation 5. The estimated regression coefficients for price impact are very similar to those obtained in the original model. For example, using the linear price impact specification, the price impact estimate, $\hat{\lambda}$, is 1.875 whereas in our original specification, the estimate is 1.811. Similarly, we find that the sets of *skilled* and *unskilled* investors from both models are very similar. For example, we find that the exact same set of 40 (resp., 46) investors are identified as *skilled* (resp. *unskilled*) in both models. Collectively, this common group nearly constitutes 90% of the original set of the identified traders suggesting that our results are robust to the choice of alpha specification.

5.3. Spread Component in Price Impact

We can generalize our price impact term by including a spread component. Due to a liquidity premium, the price impact during an execution may be higher for a less liquid stock, keeping all else equal. To test the robustness of our results with respect to a spread component, we explore an alternative specification with an additional independent variable controlling for the spread. If the model with the spread component explains much higher variation without the skill terms, accounting for short-term trading skill may lose its attractiveness. Formally, we have the following skill estimation models:

$$
\operatorname{sgn}(Q_i)\log\left(\frac{P_{i,T_i}}{P_{i,0}}\right) = \sigma_i\sqrt{T_i} \sum_{j=1}^{J} \mathbb{I}_{\{c(i)=j\}}\alpha_{c(i)} + \lambda\sigma_i\sqrt{T_i}\rho_i^{\gamma} + \delta S_i + \epsilon_i\,, \tag{6}
$$

where S_i denotes the time-weighted average bid-offer spread over the course of the execution, $\gamma = 1$ for the linear model and $\gamma = \frac{1}{2}$ for the square-root model.

[Insert Table 8 here]

Table 8 summarizes the results of these regressions along with the reduced models where we drop the skill terms. We observe that our conclusions with the original model remain unchanged. Spread parameter is only significant (at the 10% level) in the models augmented with short-term trading skill. Inclusion of the spread component do not change the R^2 in the models with and without short-term trading skill. The universe of *skilled* and *unskilled* traders and the price impact coefficient also stay largely unchanged. These results illustrate that our price impact model without a spread component provides robust findings.

5.4. Endogenous Execution Horizon

The investor has the potential to control the execution horizon by adjusting the urgency parameters of the trading algorithms. Investors having short-term horizon may be more inclined to choose faster executions compared with long-horizon investors who are satisfied with complete one-day executions. In other words, the choice of execution horizon may be endogenous.

Summary statistics in Table 2 also show that fifty percent of the executions are either less than 15 minutes or larger than 325 minutes. To address any potential bias, we consider a subset of executions for which execution horizons are very similar and do not fall into the fastest and slowest category. For this purpose, we construct another sample data where execution horizon is between 156 and 234 minutes (between 40% and 60% of a full trading day). We obtain 1651 executions submitted by 39 distinct investors.

[Insert Table 9 here]

Table 9 illustrates the regression results using this subset of data. Our earlier conclusions regarding increased explanatory power with investor identity and heterogeneity in short-term trading skill also emerge with this substantially different data set. Without accounting for short-term trading skill, R^2 values are very close to zero whereas incorporating skill terms lead to adjusted R^2 values of 16% in both price impact specifications which is again an order of magnitude difference. We also observe pronounced heterogeneity in short-term trading skill with roughly 50% of the investors being either *skilled* or *unskilled*. Finally, we note that in both models, price impact coefficients differ significantly when short-term trading skill is taken into account. Our findings suggest that earlier findings from the original data set remain unchanged qualitatively when using the restricted data set, in which execution horizon is no longer a strategic choice and is taken to be roughly one half of a trading day.

5.5. Accounting for Market Returns

Our model uses raw execution returns as dependent variables as outlined in equations (2) and (3). Given that active fund managers are evaluated against market-driven benchmarks, we can adjust our model specifications with benchmarked returns. For this reason, we explore an alternative specification in a one factor asset-pricing model dealing with abnormal returns, expressed as the difference between the execution return of a single asset and the market return. Consequently, this adjusted model illustrates that short-term trading skill can also be quantified in a benchmark setting.

In the following regressions, we check whether such an abnormal return specification results in different findings with respect to significance of short-term trading skill and price impact. Formally, we have the following specifications:

$$\text{sgn}\left(Q_i\right)\log\left(\frac{P_{i,T_i}}{P_{i,0}}\right) - r_i = \sigma_i\sqrt{T_i}\sum_{j=1}^{J}\mathbb{I}_{\{c(i)=j\}}\alpha_{c(i)}^{\mathsf{mkt}} + \lambda^{\mathsf{mkt}}\sigma_i\sqrt{T_i}\rho_i + \epsilon_i, \tag{7}$$

$$\text{sgn}\left(Q_i\right)\log\left(\frac{P_{i,T_i}}{P_{i,0}}\right) - r_i = \sigma_i\sqrt{T_i}\sum_{j=1}^{J}\mathbb{I}_{\{c(i)=j\}}\alpha_{c(i)}^{\mathsf{mkt}} + \lambda^{\mathsf{mkt}}\sigma_i\sqrt{T_i}\rho_i^{1/2} + \epsilon_i, . \tag{8}$$

Here, r_i denotes the log-return of the S&P 500 index (excluding dividends) over the date of the execution, scaled proportionally with the length of the execution duration. We use the superscripts, mkt, to emphasize the difference between these and original regression models.

[Insert Table 10 here]

Table 10 summarizes the results of these regressions. When returns are benchmarked against the market return, we observe that the number of skilled (unskilled) traders decreases (increases) when compared with our original model. This result is also consistent with the traditional performance measurement studies in the mutual fund literature as in Wermers (2000). We observe that accounting for market returns does also increase the goodness of fit slightly with an adjusted R^2 of 11.8%. Our main result from this analysis is the sustained heterogeneity in the short-term trading skill. The results support our earlier findings regarding the significance of short-term trading skill and price impact.

6. Skill and Predictability of Returns

Our estimation methodology allowed us to identify skilled and unskilled traders qualified by a short-term Sharpe ratio over the interval of execution. If skilled trading arises from superior predictive ability, as is our interpretation, then skilled trading should also be predictive of *future* returns beyond the execution horizon. The construction of a profitable trading strategy based on the persistence of short-term trading skill would provide evidence for our interpretation. In this section, we propose a long-short portfolio strategy that exploits the classification of traders based on skill. Our interpretation would be supported if this trading strategy were to generate significant abnormal returns.

6.1. Data

We use the estimated short-term predictive ability of each trader from our model estimation in Section 4. In order to test our trading strategy, we complement the original data set by extracting the next day returns of the executed stock from the CRSP database. We download daily four factor data based on portfolio returns of market risk premium, HML, SMB and UMD (momentum) and the risk-free rate from Ken French's webpage[6].

6.2. Construction of Trading Strategy

We construct a simple trading strategy according to the sign of the estimated short-term trading skill for each trader. We will restrict our universe of traders to *skilled* and *unskilled* based on our earlier definition in Section 4. The main idea of the strategy is to follow the trades of the

[6]See Ken French's data library at `http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html`.

skilled investors and perform trades in the opposite direction of what unskilled investors trade. This strategy is motivated by the hypothesis that skill should be persistent in the short-run.

We rebalance our portfolio at the end of each trading day between January 3rd, 2011 and June 29th, 2012. Specifically, we construct our long-short portfolio specifically as follows:

- On trading day t, there are L_t distinct stocks that skilled traders have bought and unskilled traders have sold. Then, at the beginning of $(t+1)$th trading day, our long portfolio will put equal (positive) weight to each of these securities with a maximum possible weight of w_{max}. Formally, each security will have a weight of $\min\left(\frac{1}{L_t}, w_{\mathsf{max}}\right)$.

- On trading day t, there are S_t distinct stocks that skilled traders have sold and unskilled traders have bought. Then, at the beginning of $(t+1)$th trading day, our short portfolio will put equal (negative) weight to each of these securities with a minimum possible weight of $-w_{\mathsf{max}}$. Formally, each security will have a weight of $\max\left(-\frac{1}{S_t}, -w_{\mathsf{max}}\right)$.

- The remaining portfolio is invested in the risk-free rate. Formally, at the beginning of $(t+1)$th trading day, the weight in the risk-free security is

$$w_{\mathsf{rf},t+1} = 1 - L_t \min\left(\frac{1}{L_t}, w_{\mathsf{max}}\right) - S_t \max\left(-\frac{1}{S_t}, -w_{\mathsf{max}}\right).$$

We set w_{max} to 5% as a base calibration and also provide sensitive analysis for this parameter.

6.3. Results

Using this long-short trading strategy, we obtain the time-series of returns for our strategy resulting from 377 trading days. The average daily return from our trading strategy is 6.91 bps corresponding to 17.4% of annualized return. The standard deviation of the daily return resulting from our strategy is 71.23 bps corresponding to annualized volatility of 11.3%. The resulting annualized Sharpe ratio is 1.94. In order to ensure that these raw return statistics are abnormal, we regress our strategy returns against the four-factor model of Carhart (1997).

[Insert Table 11 here]

Table 11 shows that our long-short portfolio returns have statistically significant Jensen's alpha. In annualized terms, employing our strategy earns an excess return of 18.8% which is also economically significant. These abnormal return statistics provide further supporting evidence of our classification of traders with regards to short-term predictive ability. Note that this also illustrates that the trades of the skilled and unskilled investors do predict the returns of the next day, which is actually a stronger result.

Table 11 also illustrates that our results are robust to the choice of w_{max}. We use three other choices of w_{max} with 2.5%, 10% and 15% and in each scenario we obtain statistically significant excess returns for our strategy. The annualized excess returns are between 13.9% and 20.3% and Sharpe ratios are between 1.54 and 1.99.

7. Cross-Sectional Variation in Short-Term Predictive Ability

In Section 4, we observed substantial variation in the estimated short-term predictive ability for each investor. In order to understand the source of this variation, in this section, we analyze how the predictive ability varies with respect to *ex ante* and *ex post* execution characteristics. *Ex ante* execution characteristics are stock-level or market-level characteristics that are in the information set of the broker and investor at the time of the trading decision. These can include, for example, past-return data, volatility or historical likelihood of a certain order size. On the other hand, *ex post* execution characteristics are only available after the end of the order execution. Execution duration, participation rate, or any interval statistic regarding the execution are examples of *ex post* execution characteristics.

For both *ex ante* and *ex post* analysis, we only consider the short-term predictive ability estimates from the linear price impact model as specified in Equation (2). Using these estimates, we associate each execution i with its investor's skill estimate, $\hat{\alpha}_{c(i)}$. In other words, each execution carries the skill estimate of the executing investor. We use decile analysis to study how short-term predictive ability varies with various characteristics. On each trading day, we distribute all executions into deciles by sorting on *ex ante* or *ex post* execution characteristic. The corresponding short-term trading skill estimate for each decile is the average of all estimates in that decile. With this approach, for each execution characteristic and every decile, we obtain a time-series data from 377 trading days. This choice of data construction will allow us to estimate decile averages across time and their corresponding standard errors. We will adjust the standard errors for heteroscedasticity and auto-correlation throughout the analysis.

7.1. *Ex Ante* Execution Characteristics

If we can identify any *ex ante* execution characteristic that can explain the performance differences between skilled and unskilled short-term trading, we can construct a proxy for short-term predictive ability with high accuracy exploiting these variables. For example, this proxy would be very valuable for a broker quoting an *ex ante* cost estimate of the order submitted by an anonymous investor. These predictive variables would also be very useful in order to uncover the source of short-term trading skill. This analysis will potentially reveal how the skilled investors are able to predict short-term returns systematically.

We examine the variation on predictive ability using six *ex ante* execution characteristics.

- **Short, Medium and Long-Term Momentum.** The first three measures are based on possible trend-following strategies based on past returns. We conjecture that skilled short-term investors may use past return data in various horizons to make trading decisions. We consider short, medium and long-term momentum measures defined by the cumulative return over the past week, month and 12 months, respectively. In order to compare across assets with various volatility levels, we normalize the returns by the average volatility of the asset over the past month. We construct deciles on these characteristics by multiplying the intended trade direction ($\text{sgn}(Q_i)$) with the normalized short-term or long-term past-return. In other words, we conjecture that skilled investors want to buy (resp, sell) an asset if its historical return is higher (resp., lower).

- **Normalized Order Quantity.** As a fourth measure, we consider the ratio between the submitted order quantity and the average traded volume over the past month. Here, we hypothesize that as the ratio gets larger, it may signify more confidence over the trade to be executed.

- **Average Spread and Volatility.** Our final two measures test whether predictive ability stems from opportunistic risk-bearing during periods of increased uncertainty. One potential source of skill could originate from making a decision to trade when there is a pronounced dispersion about the fundamental value of the asset. Here, we interpret average percentage bid-offer spread and volatility over the past month as (imperfect) measures of market liquidity and uncertainty.

[Insert Figure 6 and Table 12 here]

Figure 6 illustrates decile averages and standard error bars for all of the six *ex ante* execution characteristics. We observe consistently increasing deciles for short and medium-term momentum and executed quantity as a fraction of past average daily volume over the past month. We do not obtain similar trends for the remaining three characteristics. Since our skill measure is risk-adjusted, it is not surprising not to see any consistent pattern with regards to the opportunistic risk-bearing hypothesis. Table 12 reports these statistics formally and tests whether there is any statistical significance for the difference in means between High and Low deciles. We obtain high significance and economically large differences in means for short and medium term momentum and the ratio of number of shares to be executed to its past-month average. We observe that the difference in means for the short-term momentum is considerably higher than that of medium-term.

Our results provide evidence for short-term trading skill emerging from chasing predictability from analyzing short-term price runs. Unskilled investors in our sample do not follow momentum strategy as consistent with the broader literature on diverse strategies of institutional trading (e.g., Lakonishok et al. (1992)). Finally, we find that another reflection of skilled trading is the large trade size of the order relative to its past average daily volume of the asset.

7.2. *Ex Post* Execution Characteristics

In addition to the analysis of predictive characteristics of short-term trading skill, in this section, we examine how *ex post* execution characteristics co-vary with the level of short-term predictive ability.

We analyze the co-variation with predictive ability using five *ex post* execution characteristics.

- **Execution Duration.** The investor specifies the urgency of his execution adjusting the parameters of the broker's trading algorithm and the resulting execution duration is affected by this initial choice. Due to the trade-off between incurring cost from aggressive trading versus the risk of losing the predicted excess return, it is interesting to uncover the duration choice of the skilled investors.

- **Normalized Order Quantity.** We test two measures related to the executed quantity. We use participation rate and the ratio of executed quantity to daily volume on the day of execution. These size measures also signify the same trade-off we noted with duration. The investor can exploit larger profit with higher size but will be exposed to higher execution cost.

- **Average Spread and Volatility.** Our final two measures examine whether the presence of predictive ability is associated with lower liquidity or greater uncertainty. We use average proportional bid-offer spread and volatility over the course of the execution as measures of liquidity and uncertainty.

[Insert Figure 7 and Table 13 here]

Figure 7 illustrates the average (across time) short-term predictive ability estimates for each characteristic decile along with its corresponding standard error. We observe that there is a consistent trend between the deciles of skill estimate and four *ex post* execution characteristics. Average skill is positively correlated with shorter executions and larger sized executions, measured by the fraction of daily volume and participation rate, and bid-offer spread. Interval volatility does not exhibit a pronounced correlation with average skill. Table 13 reports decile statistics formally and test whether there is any statistical significance for the difference in means between High and Low deciles. We obtain high significance for execution duration, normalized quantity measures and bid-offer spread. These differences are also economically large especially for duration and normalized quantity measures. For example, in annualized terms, the fastest completed executions have a short-term Sharpe ratio difference of 0.73 compared to the slowest executed trades. Similarly, for the highest decile sorted by participation rate (respectively, the highest fraction of daily volume), average Sharpe ratio is 0.96 (resp., 0.56) higher than that of the lowest decile.

These results are economically easy to interpret. For a skilled investor, it makes more sense to execute his view on the asset return by trading fast. As we also detected in the previous section, a skilled investor is also more likely to trade in large size, measured either as a fraction of daily volume or as a participation rate. Our results on interval bid-ask spread are also consistent with the larger literature on information asymmetry. If we consider short-term skill estimates as a measure of information asymmetry, we observe that higher spread deciles coincide with higher information asymmetry.

We also repeat our analysis using double sorts on execution characteristics. This analysis illustrates the variation in short-term alpha term when two endogenous execution characteristics are varied. We observe that the monotonicity results in our single sort analysis for execution duration, participation rate and daily fraction of executed quantity largely remain the same.

8. Implications for Execution Quality

Our estimation results illustrated that the price impact measures are biased if investor's short-term trading skill is ignored. Our model implies that traditional measures of execution cost will also suffer from the same bias in the presence of systematic short-term predictive ability. Controlling for execution characteristics, our model predicts that execution costs of *skilled* (resp.,*unskilled*) short-term investors will be higher (resp.,lower).

In order to explore this hypothesis, we use implementation shortfall (IS) as a measure of execution cost as introduced by Perold (1988). IS is the widely preferred measure of trading cost for institutions and has been frequently employed in the literature to proxy institutional trading cost. It is computed as the normalized difference between the average execution price and the price of the asset prior to the start of the execution. Formally, the IS of ith execution in our data is given by

$$\mathrm{IS}_i = \operatorname{sgn}(Q_i)\,\frac{P_i^{\mathrm{avg}} - P_{i,0}}{P_{i,0}},$$

where P_i^{avg} is the volume-weighted execution price of the parent order.

8.1. Known Investor Identities

We use in-sample and out-of-sample data from Section 5.1. Using our short-term skill estimates from the in-sample regression, we create a dummy variable for skilled and unskilled investors as follows. Using our earlier definitions for *skilled* and *unskilled* investors from Section 4, IsSkilled_i (resp., IsUnskilled_i) takes a value of 1 if the ith execution is sent by a *skilled* (resp., an *unskilled*) investor.

[Insert Table 14 here]

In Table 14, using the out-of-sample data, we regress implementation shortfall on our skill dummies and execution level control variables including participation rate, interval bid-offer spread, interval mid-quote volatility, execution duration, and average bid-offer spread and mid-quote volatility over the past month. We are again concerned with stock-level and calendar day-level residual correlation and correct our standard errors by clustering at the stock and calendar day level.

We observe that implementation shortfall is highly positively correlated with being *skilled* investor, participation rate and interval proportional bid-offer spread and is highly negatively correlated with being *unskilled* investor. These results illustrate that heterogeneity in short-term predictive ability has a major impact on the quoted execution cost. The cost difference between *skilled* and *unskilled* investors is approximately 25 bps which is also economically significant as the median implementation shortfall in our out-of-sample data is only 5 bps.

This result has important implications for execution quality in various market venues including dark markets. Obtaining high execution quality in a particular venue does not immediately imply that the venue has superior execution quality. Given the heterogeneity in short-term predictive ability, our results suggest that without conditioning on investor type, mere comparison of execution costs cannot be conclusive. A certain market can be heavily used by skilled (resp., unskilled) investors and this may lead to an understatement (resp., overstatement) of execution quality [7]. Our model can be utilized to correct for this bias by accounting for the heterogeneity in short-term trading skill.

[7] For example, Zhu (2014) and Iyer et al. (2015) find that informed traders strategically choose the lit markets over dark pools for their execution needs.

8.2. Unknown Investor Identities

In Section 7.1, we found out that short-term momentum and quantity-to-be-traded as expressed as a fraction of average daily volume over the past month are positively correlated with short-term trading skill. We can exploit these characteristics and use them as signals to proxy for short-term trading skill in the absence of investor identities.

We again use the out-of-sample data in Section 5.1. We will provide an implementable strategy by only using the data available in the broker's information set. For every execution on trading day t, we first compute its "momentum" signal by multiplying the intended trade direction with the past week return of the stock normalized by its average monthly volatility. Second we compute its "size" signal by dividing the incoming order size by the stock's average daily volume over the past month. We use quantiles to measure whether an incoming order has the highest momentum or size signal. We use historical execution data available on the date of the execution to assign each order to its appropriate bin. In other words, using all the available executions before day t, we compute quantile breakpoints for these two signals. Using these breakpoints, we find the corresponding quantile for every execution submitted on day t. These assignments cumulatively form momentum and size quantiles. We then move on labeling the executions on the day $(t+1)$ by expanding our quantile breakpoints including the just-labeled executions from day t. All of the executions on the first day of our out-of-sample data constitute the base for our initial quantile breakpoints and we start labeling executions from the next trading day with the updated quantile breakpoints from the previous trading day.

We predict that an incoming order is sent by a *skilled* investor if the execution is in the top (fifth) quantile for both momentum and size quantiles. Similarly, we conjecture that the order is sent by an *unskilled* investor if the execution is in the bottom (first) quantile for both momentum and size quantiles. We create the dummy variable IsSkilled_i (resp., IsUnskilled_i) that takes a value of 1 if the ith execution is in the top (resp, bottom) quantiles of momentum and size.

In Table 15, we regress implementation shortfall on our skill dummies and execution level control variables including participation rate, interval bid-offer spread, interval mid-quote volatility, execution duration, and average bid-offer spread and mid-quote volatility over the past month.

We observe that the same variables – skilled dummies, participation rate, and interval spread – are again significant with the same signs as in the previous section. These results illustrate that proxies for short-term predictive ability have also a major impact on execution cost. The cost difference between skilled and unskilled investors is approximately 13 bps which is also economically significant. Compared to our identification with investor-specific ability estimates, our proxies were able to pick up more than 50% of the cost differential. These results support more evidence of the significant impact of short-term predictive ability on quoted execution costs and underscores that investor identities are not even needed in order to estimate execution costs *ex ante* for potentially skilled or unskilled investor.

9. Conclusion

Statistical models for short-term returns observed during the execution of a large order typically have low explanatory power. It is difficult to separate price impact due to demand for liquidity from predicting price changes due to trading skill. Consequently, trading cost models are estimated on large order samples, where the effect of short-term predictive ability is expected to cancel out.

In this paper, we first propose a model to explain the variation in short-term returns with the short-term trading skill of the investors along with a parametric modeling of price impact. Motivated by the performance metrics for the fund management industry, we measure trading skill in a risk-adjusted way using short-term Sharpe ratios. We estimated our model on a large sample of executions with masked investor identities and our results show that incorporating short-term predictive ability offers drastic improvements in explaining the variation in security returns over an execution horizon. We also observe that ignoring short-term trading skill may lead to biased price impact estimates which are economically large.

The estimated trading skill is specific to each investor in the sample, and can be used to classify investors according to the success of their predictive ability. We find that in addition to the presence of skilled investors, a significant portion of the investor universe is unskilled in the sense that their trades are in the opposite direction of future short-term price movements. This cross-sectional variation implies a pronounced heterogeneity in short-term trading skill among institutional investors. This classification is robust and has predictive power about the future trading performance of an investor. In order to test for persistence in the short-term, we propose a trading strategy exploiting the signs of estimated skill coefficients and find that this strategy generates economically substantial abnormal returns even against a benchmark Fama-French four-factor model.

We find strong evidence that skilled (unskilled) traders are able to predict short-term returns by following a momentum (contrarian) strategy in the short-term. Skilled traders reveal their confidence by trading in large amounts relative to the historical volume observed in the stock. We propose these two measures, short-term momentum and relative size, in order to predict whether an order is submitted by a skilled or unskilled investor. Furthermore, we find that the magnitude of the investor's trading skill can also be explained by *ex post* order and microstructure variables. On average, higher trading skill tends to emerge in executions with shorter durations, larger sizes and higher bid-offer spreads.

Besides a better explanation of short-term returns, short-term trading skill has major implications on the analysis of execution quality. Our results illustrate that all else equal, a venue populated with unskilled traders would report better execution quality statistics compared with a venue populated with skilled traders. Our paper proposes a simple methodology to identify these skilled and unskilled investors and enable policymakers to access better measures of execution quality. Moreover, our findings have several other practical applications. For example, an agency broker can use the historically estimated trader skill to advise individual investors on the choice of algorithm and parameter settings. Tracking the predictive ability of an investor through time gives a measure of trading efficiency, that would be of interest to this investor. Such measures can complement the traditional transaction cost analysis (TCA) that brokers typically provide. Our proxies for short-term trading skill can also be utilized when there is limited data on the historical executions of a particular investor.

References

R. Almgren, C. Thum, E. Hauptmann, and H. Li. Direct estimation of equity market impact. *Risk*, 18(7):58–62, 2005.

Robert Almgren. Execution costs. *Encyclopedia of quantitative finance*, 2008.

Amber Anand, Paul Irvine, Andy Puckett, and Kumar Venkataraman. Performance of institutional trading desks: An analysis of persistence in trading costs. *Review of Financial Studies*, 25(2): 557–598, 2012.

Clifford S Asness, Tobias J Moskowitz, and Lasse Heje Pedersen. Value and momentum everywhere. *The Journal of Finance*, 68(3):929–985, 2013.

M. Baron, J. Brogaard, and A Kirilenko. The trading profits of high frequency traders. Working paper, 2013.

Dimitris Bertsimas and Andrew W Lo. Optimal control of execution costs. *Journal of Financial Markets*, 1(1):1–50, 1998.

Hendrik Bessembinder and Herbert M Kaufman. A comparison of trade execution costs for nyse and nasdaq-listed stocks. *Journal of Financial and Quantitative Analysis*, 32(03):287–310, 1997a.

Hendrik Bessembinder and Herbert M Kaufman. A cross-exchange comparison of execution costs and information flow for nyse-listed stocks. *Journal of Financial Economics*, 46(3):293–319, 1997b.

J.-P. Bouchaud, J. D. Farmer, and F. Lillo. How markets slowly digest changes in supply and demand. In *Handbook of Financial Markets: Dynamics and Evolution*, pages 57–156. Elsevier: Academic Press, 2008.

Jonathan Brogaard, Terrence Hendershott, Stefan Hunt, and Carla Ysusi. High-frequency trading and the execution costs of institutional investors. *Financial Review*, 49(2):345–369, 2014.

Mark M Carhart. On persistence in mutual fund performance. *The Journal of finance*, 52(1):57–82, 1997.

George C. Chacko, Jakub W. Jurek, and Erik Stafford. The price of immediacy. *The Journal of Finance*, 63(3):1253–1290, 2008.

Louis KC Chan and Josef Lakonishok. The behavior of stock prices around institutional trades. *The Journal of Finance*, 50(4):1147–1174, 1995.

Adam D Clark-Joseph. Exploratory trading. *Unpublished working paper). Harvard University, Cambridge MA*, 2013.

Randolph B Cohen, Joshua D Coval, and L'uboš Pástor. Judging fund managers by the company they keep. *The Journal of Finance*, 60(3):1057–1096, 2005.

Pierre Collin-Dufresne and Vyacheslav Fos. Do prices reveal the presence of informed trading? *Journal of Finance*, 70(4):1555–1582, 2015.

Karl B Diether, Kuan-Hui Lee, and Ingrid M Werner. Short-sale strategies and return predictability. *Review of Financial Studies*, 22(2):575–607, 2009.

D. Easley, M. López de Prado, and M. O'Hara. Flow toxicity and liquidity in a high frequency world. *Review of Financial Studies*, 25(5):1457–1493, 2012.

L. R. Glosten and P. R. Milgrom. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1):71–100, 1985.

Paul A Gompers and Andrew Metrick. Institutional investors and equity prices. *Quarterly Journal of Economics*, pages 229–259, 2001.

J. Hasbrouck. *Empirical Market Microstructure*. Oxford University Press, 2007.

Terrence Hendershott, Dmitry Livdan, and Norman Schürhoff. Are institutions informed about news? *Journal of Financial Economics*, 117(2):249–287, 2015.

Roger D. Huang and Hans R. Stoll. Dealer versus auction markets: A paired comparison of execution costs on NASDAQ and the NYSE. *Journal of Financial Economics*, 41(3):313 – 357, 1996.

K. Iyer, R. Johari, and C. C. Moallemi. Welfare analysis of dark pools. SSRN Working paper, 2015.

Michael Johannes, Arthur Korteweg, and Nicholas Polson. Sequential learning, predictability, and optimal portfolio returns. *The Journal of Finance*, 69(2):611–644, 2014.

Charles M Jones. What do we know about high-frequency trading? *Columbia Business School Research Paper*, (13-11), 2013.

Marcin Kacperczyk, Clemens Sialm, and Lu Zheng. On the industry concentration of actively managed equity mutual funds. *The Journal of Finance*, 60(4):1983–2011, 2005.

Ron Kaniel and Hong Liu. So what orders do informed traders use?*. *The Journal of Business*, 79 (4):1867–1913, 2006.

Donald B Keim and Ananth Madhavan. Transactions costs and investment style: An inter-exchange analysis of institutional equity trades. *Journal of Financial Economics*, 46(3):265–292, 1997.

Donald Bruce Keim and Ananth Madhavan. The upstairs market for large-block transactions: Analysis and measurement of price effects. *Review of Financial Studies*, 9(1):1–36, 1996.

A. P. Kyle and A. Obizhaeva. Market microstructure invariants: Theory and empirical tests. SSRN Working paper, 2014.

A. S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335, November 1985.

Josef Lakonishok, Andrei Shleifer, and Robert W Vishny. The impact of institutional trading on stock prices. *Journal of financial economics*, 32(1):23–43, 1992.

C. Maglaras, C. C. Moallemi, and H. Zheng. Queueing dynamics and state space collapse in fragmented limit order book markets. Working paper, 2012.

Harry Mamaysky, Matthew Spiegel, and Hong Zhang. Estimating the dynamics of mutual fund alphas and betas. *Review of Financial Studies*, 21(1):233–264, 2008.

Tobias J Moskowitz, Yao Hua Ooi, and Lasse Heje Pedersen. Time series momentum. *Journal of Financial Economics*, 104(2):228–250, 2012.

A. Obizhaeva. Liquidity estimates and selection bias. Working paper, 2009.

Maureen O'Hara and Mao Ye. Is market fragmentation harming market quality? *Journal of Financial Economics*, 100(3):459 – 474, 2011. ISSN 0304-405X. doi: http://dx.doi.org/10.1016/j.jfineco.2011.02.006.

Andre F Perold. The implementation shortfall: Paper versus reality. *The Journal of Portfolio Management*, 14(3):4–9, 1988.

Mitchell A Petersen. Estimating standard errors in finance panel data sets: Comparing approaches. *Review of financial studies*, 22(1):435–480, 2009.

Andy Puckett and Xuemin Sterling Yan. The interim trading skills of institutional investors. *The Journal of Finance*, 66(2):601–633, 2011.

Lin Tong. A blessing or a curse? the impact of high frequency trading on institutional investors. *SSRN Working Paper*, 2015.

N Torre and Mark J Ferrari. The market impact model. *Horizons, The Barra Newsletter*, 165, 1998.

Russ Wermers. Mutual fund performance: An empirical decomposition into stock-picking talent, style, transactions costs, and expenses. *The Journal of Finance*, 55(4):1655–1703, 2000.

Xuemin (Sterling) Yan and Zhe Zhang. Institutional investors and equity returns: Are short-term institutions better informed? *Review of Financial Studies*, 22(2):893–924, 2009.

Haoxiang Zhu. Do dark pools harm price discovery? *Review of Financial Studies*, 27(3):747–789, 2014.

Figure 1: An illustration of the price change attribution of an asset during an investor's sale through a broker execution.

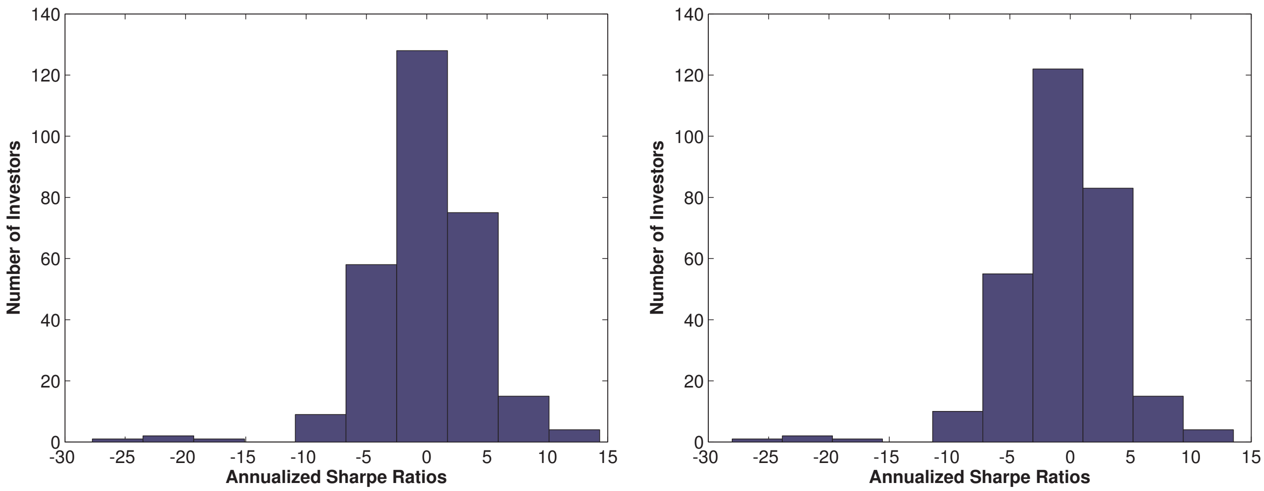

Figure 2: Histogram of all investor skill estimates expressed as annualized Sharpe ratios, when the price impact is proportional to the participation rate (left) and when the price impact is proportional to the square root of participation rate (right).

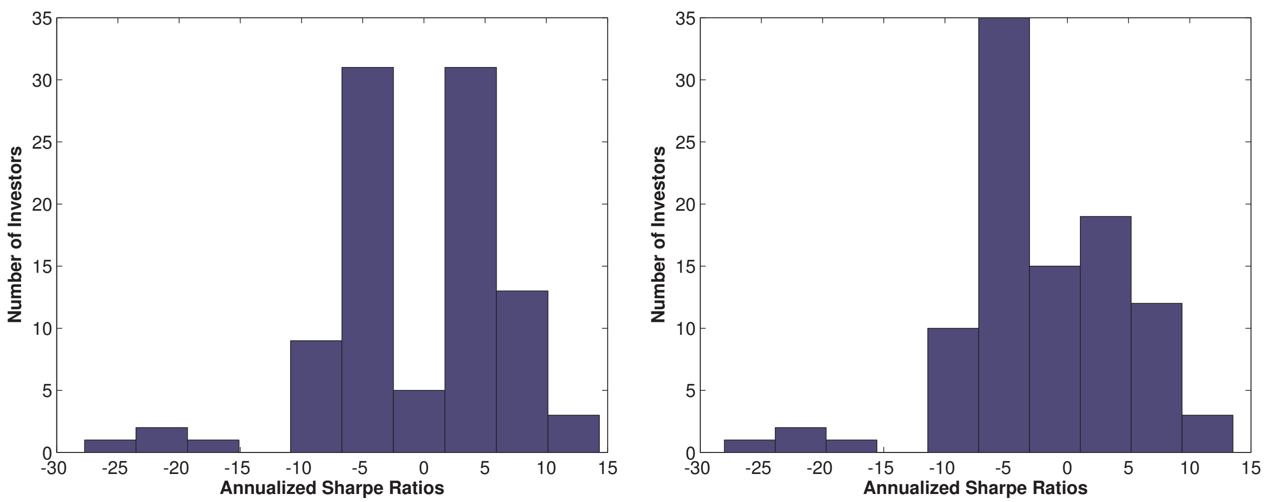

Figure 3: Histogram of statistically significant investor skill estimates, expressed as annualized Sharpe ratios, when price impact is proportional to the participation rate (left) and when the price impact is proportional to the square root of the participation rate (right).

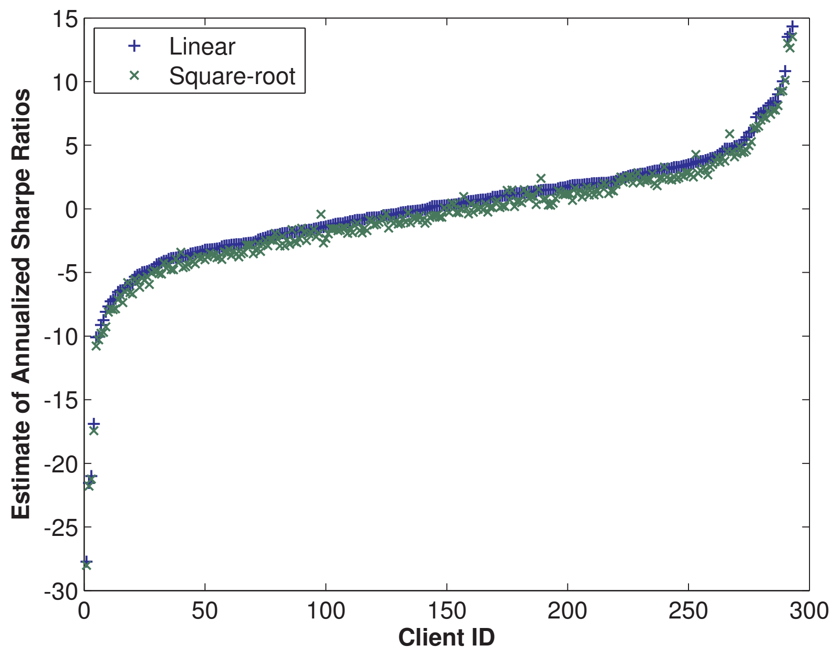

Figure 4: The difference in trading skill estimates between the linear and square root price impact models in Equations (2) and (3).

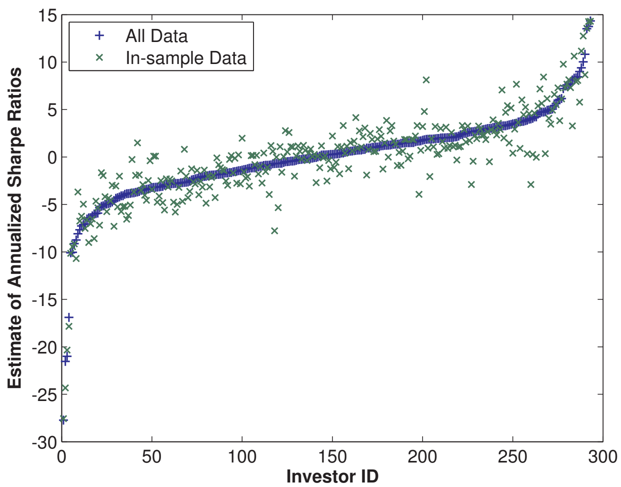

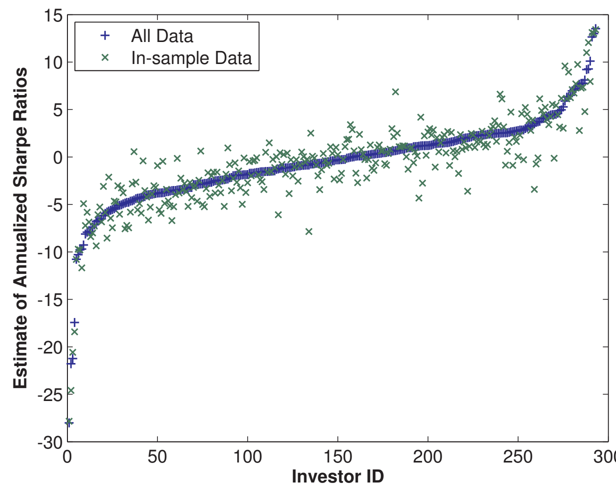

Figure 5: As a robustness check, we compare trading skill estimates computed from the complete data set and a randomly constructed in-sample data set in which for every investor only random half of his executions are considered. We use two price impact specifications for our comparisons: the linear model (left) and the square-root model (right).

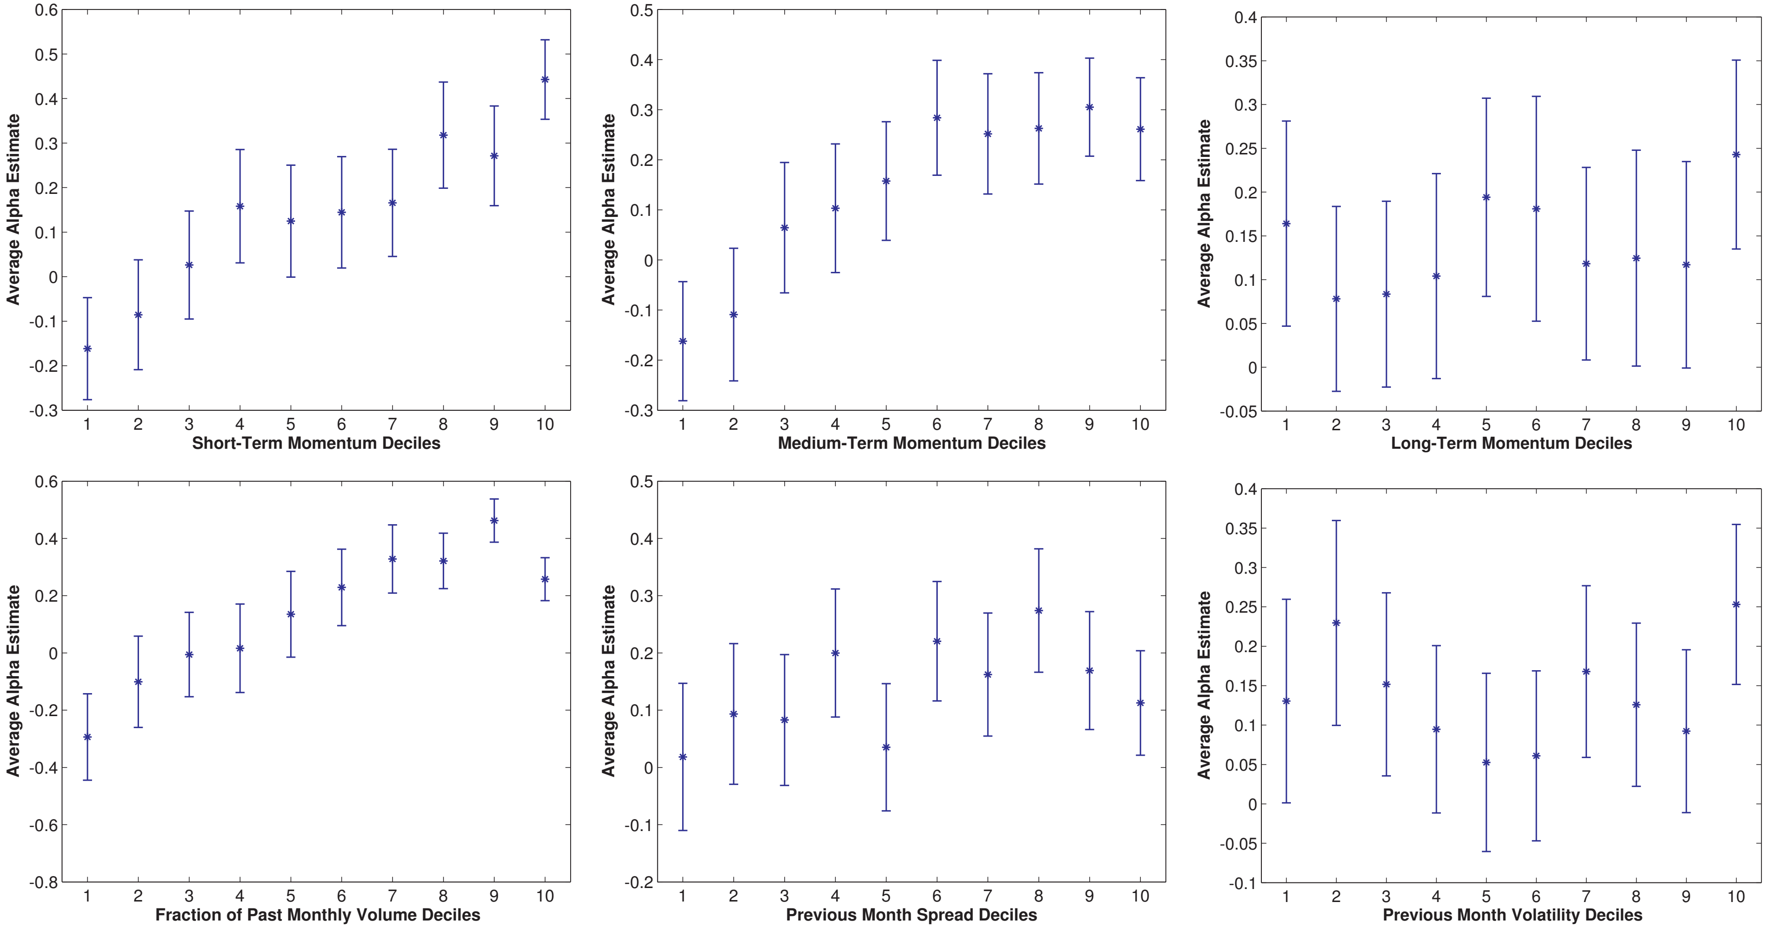

Figure 6: This figure illustrates annualized average short-term predictive ability estimates for each decile from a single sort on *ex ante* execution characteristics including volatility normalized short-term (past one-week raw return) (top-left), medium-term (past one-month raw return) (top-middle), long-term momentum (past 12-month raw return) (top-right), size of the order as a fraction of daily average volume over the past month (bottom-left), daily average bid-offer spread over the past month (bottom-middle) and daily average mid-quote volatility over the past month (bottom-right). For every trading day in our sample, we sort executions on these characteristics and and report the average of corresponding trading skill estimates in each decile. Error bars report heteroscedasticity and auto-correlation consistent standard errors.

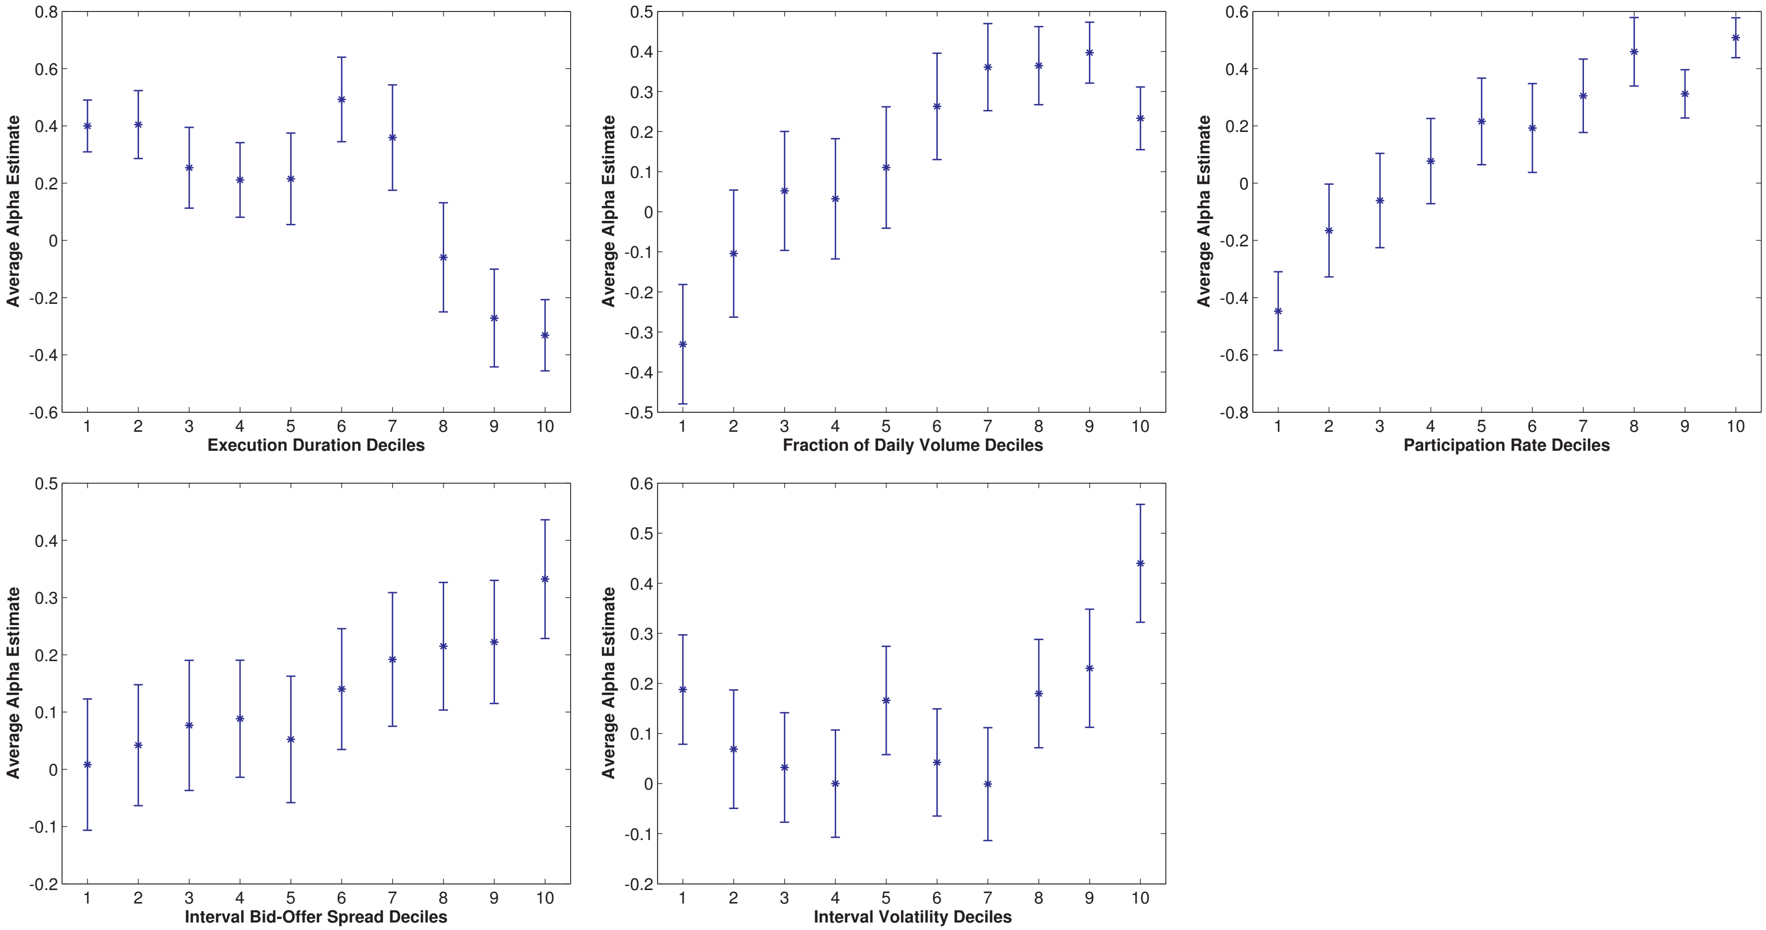

Figure 7: This figure illustrates annualized average short-term predictive ability estimates for each decile from a single sort on contemporaneous execution characteristics including execution duration (top-left), executed quantity as fraction of daily volume (top-middle), participation rate (top-right), interval bid-offer spread (bottom-left) and mid-quote volatility (bottom-middle). For every trading day in our sample, we sort executions on these characteristics and report the average annualized trading skill estimates in each decile. Error bars report heteroscedasticity and auto-correlation consistent standard errors.

Statistic	Interval Return (%)	Bid Ask Spread (bps)	Average Daily Volatility (%)	Execution Duration (mins)	Participation Rate (%)	Percentage of Daily Volume (%)
Min.	-18.102	0.68	0.128	5.00	0.0001	<0.01
1st Qu.	-0.307	2.45	0.980	15.23	0.17	0.06
Median	0.000	3.35	1.267	59.82	1.59	0.26
Mean	0.003	4.16	1.417	148.13	6.44	0.64
3rd Qu.	0.333	4.87	1.685	324.50	10.50	0.72
Max.	10.793	54.76	25.740	390.00	100.00	28.25

Table 1: Summary statistics for the main attributes in our execution data. The bid ask spread is normalized using the mid-quote price. Average daily volatilities are computed using the previous 20 trading days before the execution date. The duration of a full trading day in U.S. equity markets is 390 minutes.

	Linear		Square root	
Trading Skill?	Yes	No	Yes	No
Intercept (bps)	1.26	−1.34	1.51	−3.27
	(1.24)	(2.42)	(1.22)	(2.69)
λ	1.811***	2.277***	0.744***	0.837***
	(0.291)	(0.363)	(0.173)	(0.154)
Number of *skilled* investors	49	N/A	35	N/A
Number of *unskilled* investors	48	N/A	63	N/A
R^2	10.5%	0.5%	10.5%	0.4%
Adj. R^2	10.1%	0.5%	10.0%	0.4%

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Table 2: Regression results for two price impact models with and without trading skill terms. Two price impact specifications are estimated, linear and square root. We label an investor as *skilled* (resp, *unskilled*) if his short-term trading skill estimate is positive (resp., negative) and is statistically significant under 10% level. In each column, we report estimated coefficients and their standard errors, adjusted by clustering on calendar day and stock as suggested by Petersen (2009).

	Linear	Square root
$\hat{\lambda} - \hat{\lambda}^{\mathsf{base}}$	-0.467^{***}	-0.096^{***}
	(0.012)	(0.007)
N	$1,000$	$1,000$

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Table 3: Bootstrapping results for the difference in price impact coefficients by constructing $1,000$ random datasets each with $10,000$ executions. In each column, we report the mean difference between the price impact coefficients and its corresponding standard errors.

	Linear	Square root
Number of *skilled* investors	49^{***}	35^{***}
	[< 0.001]	[< 0.001]
Number of *unskilled* investors	48^{***}	63^{**}
	[0.001]	[0.016]
$\hat{\lambda}$	1.811^{***}	0.744^{***}
	[< 0.001]	[< 0.001]
N	$10,000$	$10,000$

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Table 4: Permutation test results for two price impact models by shuffling investor identifiers across executions. In each column, we report estimated coefficients from the original models and their corresponding p-values in square brackets. Empirical distribution for the parameters are obtained under the null hypothesis that investor identifiers are unrelated to log-returns realized during each execution horizon.

Trading Skill?	Linear Yes	Linear No	Square root Yes	Square root No
Intercept (bps)	1.41	−0.76	1.55	−2.93
	(1.54)	(2.41)	(1.53)	(2.62)
λ	1.793***	2.160***	0.767***	0.838***
	(0.310)	(0.366)	(0.179)	(0.156)
Number of *skilled* investors	42	N/A	32	N/A
Number of *unskilled* investors	45	N/A	53	N/A
R^2	11.4%	0.5%	11.4%	0.5%
Adj. R^2	10.6%	0.5%	10.6%	0.5%

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$

Table 5: Regression results for two price impact models with an in-sample data set constructed using random half of each investor's executions. We label an investor as *skilled* (resp, *unskilled*) if his short-term trading skill estimate is positive (resp., negative) and is statistically significant under 10% level. We estimate two price impact models with and without predictive ability of the investor. In each column, we report estimated coefficients and their standard errors, adjusted by clustering on calendar day and stock as suggested by Petersen (2009).

Model	rMSPE In-Sample	rMSPE Out-of-Sample	R^2 In-Sample	R^2 Out-of-Sample
Linear	93.19	95.06	10.6%	8.2%
Square root	93.19	95.07	10.6%	8.1%

Table 6: This table reports root mean squared prediction errors (rMSPE) between in-sample and out-of-sample execution returns and in-sample and out-of-sample R^2. Predicted execution returns use the skill and price impact coefficients estimated from in-sample data set. rMSPE values are reported in basis points.

	Linear	Square root
Intercept (bps)	1.48	1.08
	(1.03)	(1.20)
λ^{mkt}	1.875***	0.759**
	(0.243)	(0.140)
Number of *skilled* investors	45	35
Number of *unskilled* investors	61	71
R^2	10.9%	10.9%
Adj. R^2	10.5%	10.5%

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Table 7: Regression results for two price impact models using constant daily Sharpe ratio for each investor. We only consider the presence of trading skill terms. We label an investor as *skilled* (resp, *unskilled*) if his short-term trading skill estimate is positive (resp., negative) and is statistically significant under 10% level. In each column, we report estimated coefficients and their standard errors, adjusted by clustering on calendar day and stock as suggested by Petersen (2009).

	Linear		Square root	
Trading Skill?	Yes	No	Yes	No
Intercept (bps)	0.41	−1.12	0.63	−2.78
	(1.31)	(1.80)	(1.29)	(2.03)
λ	1.758***	2.284***	0.721***	0.846***
	(0.310)	(0.366)	(0.179)	(0.176)
δ (bps)	0.304*	−0.055	0.316*	−0.135
	(0.179)	(0.433)	(0.179)	(0.180)
Number of *skilled* investors	46	N/A	32	N/A
Number of *unskilled* investors	48	N/A	63	N/A
R^2	10.5%	0.5%	10.5%	0.4%
Adj. R^2	10.1%	0.5%	10.0%	0.4%

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$

Table 8: Regression results for two price impact models with an additional spread component in the price impact model. We label an investor as *skilled* (resp, *unskilled*) if his short-term trading skill estimate is positive (resp., negative) and is statistically significant under 10% level. We estimate two price impact models with and without predictive ability of the investor. In each column, we report estimated coefficients and their standard errors, adjusted by clustering on calendar day and stock as suggested by Petersen (2009).

	Linear		Square root	
Trading Skill?	Yes	No	Yes	No
Intercept (bps)	−2.29	−3.79	−3.28	−4.70
	(8.77)	(6.28)	(8.65)	(3.29)
λ	2.726**	2.455*	1.063*	0.464
	(1.243)	(1.411)	(0.554)	(0.649)
Number of executions	1651	1651	1651	1651
Number of investors	39	39	39	39
Number of *skilled* investors	9	N/A	8	N/A
Number of *unskilled* investors	9	N/A	10	N/A
R^2	17.9%	0.3%	18.0%	0.1%
Adj. R^2	15.8%	0.2%	15.9%	0.1%

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Table 9: Regression results for two price impact models for a subset of executions that lasted between 156 and 234 minutes (between 40% and 60% of a full trading day). We label an investor as *skilled* (resp, *unskilled*) if his short-term trading skill estimate is positive (resp., negative) and is statistically significant under 10% level. We estimate two price impact models with and without predictive ability of the investor. In each column, we report estimated coefficients and their standard errors, adjusted by clustering on calendar day and stock as suggested by Petersen (2009).

	Linear	Square root
Intercept (bps)	1.89	2.46
	(2.49)	(2.44)
λ^{mkt}	1.609***	0.555**
	(0.406)	(0.265)
Number of *skilled* investors	33	25
Number of *unskilled* investors	55	61
R^2	12.2%	12.2%
Adj. R^2	11.8%	11.8%

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Table 10: Regression results for two price impact models using abnormal (excess) returns. We only consider the presence of trading skill terms. We label an investor as *skilled* (resp, *unskilled*) if his short-term trading skill estimate is positive (resp., negative) and is statistically significant under 10% level. In each column, we report estimated coefficients and their standard errors, adjusted by clustering on calendar day and stock as suggested by Petersen (2009).

	(1) $w_{\mathsf{max}} = 2.5\%$	(2) $w_{\mathsf{max}} = 5\%$	(3) $w_{\mathsf{max}} = 10\%$	(4) $w_{\mathsf{max}} = 15\%$
Intercept (bps)	5.50*	7.48**	8.07**	6.59*
	(3.05)	(3.34)	(3.51)	(3.51)
Mkt-RF	−0.024	−0.044	−0.033	−0.024
	(0.059)	(0.059)	(0.037)	(0.049)
SMB	−0.090	−0.075	−0.025	−0.051
	(0.073)	(0.075)	(0.065)	(0.070)
HML	−0.080	−0.089	−0.037	0.032
	(0.091)	(0.104)	(0.115)	(0.119)
UMD	0.006	0.021	0.104*	0.114*
	(0.061)	(0.066)	(0.055)	(0.066)
Sharpe Ratio	1.61	1.94	1.99	1.54
N	377	377	377	377
R^2	0.02	0.03	0.03	0.02
Adjusted R^2	0.01	0.02	0.02	0.01

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Table 11: Regression results for long-short portfolio returns against the four factor model due to Carhart (1997). In each column, we report estimated coefficients and their standard errors, calculated using heteroscedasticity and auto-correlation consistent standard errors.

Decile	Short-Term Momentum	Medium-Term Momentum	Long-Term Momentum	Fraction of Past Monthly Volume	Past Month Bid Ask Spread	Past Month Volatility
Low	-0.162	-0.162	0.164	-0.294*	0.018	0.130
	(0.11)	(0.12)	(0.12)	(0.15)	(0.13)	(0.13)
2	-0.086	-0.109	0.078	-0.101	0.093	0.230*
	(0.12)	(0.13)	(0.11)	(0.16)	(0.12)	(0.13)
3	0.026	0.064	0.083	-0.006	0.083	0.152
	(0.12)	(0.13)	(0.11)	(0.15)	(0.11)	(0.12)
4	0.158	0.103	0.104	0.016	0.200*	0.095
	(0.13)	(0.13)	(0.12)	(0.15)	(0.11)	(0.11)
5	0.125	0.158	0.194*	0.135	0.035	0.053
	(0.13)	(0.12)	(0.11)	(0.15)	(0.11)	(0.11)
6	0.144	0.284**	0.181	0.229*	0.220**	0.061
	(0.13)	(0.11)	(0.13)	(0.13)	(0.10)	(0.11)
7	0.166	0.252**	0.118	0.328***	0.162	0.168
	(0.12)	(0.12)	(0.11)	(0.12)	(0.11)	(0.11)
8	0.318***	0.263**	0.125	0.321***	0.274**	0.126
	(0.12)	(0.11)	(0.12)	(0.10)	(0.11)	(0.10)
9	0.271**	0.305***	0.117	0.462***	0.169	0.092
	(0.11)	(0.10)	(0.12)	(0.08)	(0.10)	(0.10)
High	0.443***	0.261**	0.243**	0.258***	0.113	0.253**
	(0.09)	(0.10)	(0.11)	(0.08)	(0.09)	(0.10)
High-Low	0.605***	0.423***	0.079	0.551***	0.094	0.123
	(0.13)	(0.13)	(0.13)	(0.17)	(0.10)	(0.11)

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Table 12: This table reports annualized average short-term predictive ability estimates for each decile from a single sort on *ex ante* execution characteristics including volatility normalized short-term (past one-week raw return), medium-term (past one-month raw return), long-term momentum (past 12-month raw return), size of the order as a fraction of daily average volume over the past month, daily average bid-offer spread over the past month and daily average mid-quote volatility over the past month. For every trading day in our sample, we sort executions into deciles of these execution characteristics and report the average of corresponding alpha estimates in each decile. The row "High-Low" reports the difference in average alpha estimates between the High and Low deciles. Heteroscedasticity and auto-correlation consistent standard errors are given in parentheses.

Decile	Execution Duration	Daily Fraction of Executed Quantity	Participation Rate	Interval Bid Ask Spread	Interval Volatility
Low	0.400***	-0.331**	-0.447***	0.008	0.188*
	(0.09)	(0.15)	(0.14)	(0.11)	(0.11)
2	0.405***	-0.104	-0.166	0.042	0.069
	(0.12)	(0.16)	(0.16)	(0.11)	(0.12)
3	0.254*	0.052	-0.061	0.077	0.032
	(0.14)	(0.15)	(0.16)	(0.11)	(0.11)
4	0.211	0.032	0.077	0.088	0.000
	(0.13)	(0.15)	(0.15)	(0.10)	(0.11)
5	0.215	0.110	0.216	0.052	0.166
	(0.16)	(0.15)	(0.15)	(0.11)	(0.11)
6	0.492***	0.263**	0.192	0.140	0.042
	(0.15)	(0.13)	(0.16)	(0.11)	(0.11)
7	0.359*	0.361***	0.305**	0.192	-0.001
	(0.18)	(0.11)	(0.13)	(0.12)	(0.11)
8	-0.059	0.365***	0.459***	0.215*	0.180*
	(0.19)	(0.10)	(0.12)	(0.11)	(0.11)
9	-0.271	0.397***	0.312***	0.222**	0.230*
	(0.17)	(0.08)	(0.08)	(0.11)	(0.12)
High	-0.331***	0.233***	0.508***	0.332***	0.440***
	(0.12)	(0.08)	(0.07)	(0.10)	(0.12)
High-Low	-0.731***	0.564***	0.955***	0.324***	0.252*
	(0.14)	(0.17)	(0.15)	(0.11)	(0.13)

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Table 13: This table reports annualized average short-term predictive ability estimates from a single sort on contemporaneous execution characteristics including execution duration, executed quantity as fraction of daily volume, participation rate, interval bid ask spread and volatility. For every trading day in our sample, we sort executions into deciles of these execution characteristics and report the average of corresponding alpha estimates in each decile. The row "High-Low" reports the difference in average trading skill estimates between the High and Low deciles. Heteroscedasticity and auto-correlation consistent standard errors are given in parentheses.

	Dependent variable:
	Implementation Shortfall
IsSkilled	11.49^{***}
	(2.72)
IsUnskilled	-13.55^{***}
	(3.55)
Participation Rate	29.94^{***}
	(5.04)
Interval Spread (bps)	0.58^{**}
	(0.29)
Interval Volatility	-88.80
	(206.95)
Execution Duration	-3.43
	(3.73)
Average Daily Volatility	137.49
	(197.60)
Average Daily Spread	-0.01
	(0.13)
Constant	1.08
	(2.88)
N	31,690
R^2	0.014
Adjusted R^2	0.013

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Table 14: We regress implementation shortfall on our skill dummies and execution level control variables including participation rate, interval bid-offer spread, interval mid-quote volatility, execution duration, and average bid-offer spread and mid-quote volatility over the past month. We use the out-of-sample data constructed for robustness checks in Section 5. Using the estimation from in-sample data, IsSkilled_i (resp., IsUnskilled_i) takes a value of 1 if the ith execution is sent by a *skilled* (resp., an *unskilled*) investor (as defined in Section 4). Standard errors are given in parentheses and are adjusted by clustering on calendar day and stock as suggested by Petersen (2009).

	Dependent variable:
	Implementation Shortfall
Skilled	5.69^{***}
	(2.19)
IsUnskilled	-7.06^{*}
	(3.80)
Participation Rate	29.47^{***}
	(4.98)
Interval Spread (bps)	0.84^{**}
	(0.36)
Interval Volatility	-124.00
	(213.39)
Execution Duration	-3.30
	(3.97)
Average Daily Volatility	210.27
	(227.94)
Average Daily Spread	-0.32
	(0.33)
Constant	0.46
	(2.95)
N	31,199
R^2	0.005
Adjusted R^2	0.005

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$

Table 15: We regress implementation shortfall on our skill dummies and execution level control variables including participation rate, interval bid-offer spread, interval mid-quote volatility, execution duration, and average bid-offer spread and mid-quote volatility over the past month. We use the out-of-sample data constructed for robustness checks in Section 5. Using the proxies for short-term trading skill, IsSkilled_i (resp., IsUnskilled_i) takes a value of 1 if the ith execution is in the top (resp., bottom) quantiles of momentum and size. Standard errors are given in parentheses and are adjusted by clustering on calendar day and stock as suggested by Petersen (2009)