WARWICK

ECONOMICS

A test of speculative arbitrage: is the cross-section of volatility invariant?

Yashar H. Barardehi, Dan Bernhardt & Thomas G. Ruchti

July 2019 No: 1204

Warwick Economics Research Papers

ISSN 2059-4283 (online)
ISSN 0083-7350 (print)

A test of speculative arbitrage: is the cross-section of volatility invariant?*

Yashar H. Barardehi

Dan Bernhardt

Thomas G. Ruchti

June 27, 2019

Abstract

We derive testable implications of Kyle and Obizhaeva's (2016) notion of "bet invariance" for the cross-section of trade-time volatilities. We jointly develop theoretical foundations of "no speculative arbitrage" whose implications incorporate those of bet invariance. Our proposed test circumvents the unobservable nature of "bets." Utilizing a large sample of U.S. stocks post decimilization, we show that using realized volatilities rather than expected volatilities introduces noise that substantially biases the tests. This leads us to use estimates of normalized volatilities based on running 24 month windows. We find strong support for no speculative arbitrage at a moment in time, but not across time.

*We are grateful for comments received from Torben Andersen, Shmuel Baruch, Andy Fodor, Burton Hollifield, Tim Johnson, Steve Karolyi, Matt Shum, and seminar participants at Hull University, Ohio University, and University of Queensland. Any error is our own. Barardehi (barardehi@chapman.edu) is at the Argyros School of Business & Economics, Chapman University; Bernhardt (danber@illinois.edu) is at Department of Economics, University of Illinois and University of Warwick; Ruchti (ruchti@andrew.cmu.edu) is at Tepper School of Business, Carnegie Mellon University.

1 Introduction

Do speculators compete away abnormal profit opportunities across assets, trading conditions, and time? With free entry, speculators should be able to allocate resources to information acquisition in such a way that expected profit opportunities are equated across these dimensions. Our paper builds on existing market microstructure theories to develop novel predictions that describe how speculative arbitrage underlies systematic cross-sectional variation in intraday volatilities. We uncover strong evidence that arbitrage equalizes expected profit opportunities from speculation across assets and trading conditions at a moment in time, but not across time.

We build on the analysis of Kyle and Obizhaeva (2016; henceforth, KO) who conjecture a high-frequency efficiency result, "bet invariance." KO introduce the notion of dollar risk transfers, or the expected dollar loss or gain associated with implementing a bet, i.e., a desired change in an investor's position in a risky asset.[1] The MMI hypothesis treats bets as random variables featuring asset- and time-specific probability distributions. KO hypothesize that in a trading environment with minimal market microstructure frictions, choices of potentially-informed speculators should lead to an invariant distribution of dollar risk transfers over time and across assets. Otherwise, speculators could re-allocate resources to assets representing greater expected profit opportunities.[2]

Our paper derives and implements testable implications of bet invariance for the cross-section of volatility that rely only on publicly-available data. These implications hold under weaker identifying assumptions than those required by existing tests. We also develop an intuitive theory, "no speculative arbitrage," that produces the same predictions in this context as KO's "bet invariance." We test these implications using a large sample of liquid U.S. stocks post decimilization. Our test is unique in that it relies only on standard trading data, exploiting cross-sectional variation in volatility. We show that bet invariance is satisfied only when using expected volatility as an input, rather than realized spot volatility. We argue that is because speculators equate profit opportunities in

[1]More generally, a bet can refer to an aggregate measure of all intended positions to be taken (or left) based on the same piece of positive (negative) information.

[2]Increased algorithmic and low-latency trading has been found to be associated with improvements in measures of price efficiency and market quality. See, for example, Hendershott and Riordan (2009), Hasbrouck and Sarr (2013), O'Hara et al. (2014), Conrad et al. (2015), and Hendershott and Riordan (2009).

expectation, whereas the actual outcomes are stochastic.

Tests of invariance confront two fundamental challenges. First, invariance represents a restriction on dollar risk transfers of bets, which are intrinsically unobservable. Existing tests do not circumvent this unobservability, with the consequence that their applicability is limited to a few particular settings. For example, KO use proprietary data on trades by (non-strategic) portfolio transition specialists, whose bets they argue are plausibly observable. Our test takes observable total volatility measures as inputs, directly addressing the unobservability of bet risk transfers.

The second challenge that existing tests must confront is that using observable volatilities requires a set of strong auxiliary assumptions. In particular, they must assume that the arrival rate of bets does not vary across different public information environments, e.g, times of the trading day; and they must assume that the bet-to-total volume and volatility ratios are fixed over time and across market conditions. These auxiliary assumptions are necessary for single-asset tests because both total volume and total volatility reflect more than the arrival of speculative bets and its associated volatility.[3] We address this second challenge by expanding the implications of bet invariance to the cross-section of assets, facilitating normalization of cross-sectional variation with respect to an arbitrary benchmark asset. This normalization allows for less restrictive assumptions about variation in bet-to-total volume and volatility over time. In particular, we can accommodate common cross-stock variation over the trading day and with the level of trading activity.

Our empirical analysis focuses on normalized trade-time volatilities that (i) measure the spot volatility associated with trading a fixed-dollar value, and (ii) control for well known variation in volatility across time-of-day, as well as microstructure noise. We define a trade-time interval to be the time required for realizations of sequences of trades with a month-specific fixed-dollar value to trade: shorter time durations of trade sequences reflect more active trading conditions, i.e.,

[3] Andersen et al. (2017) identify a setting where these assumptions may be reasonable. They focus on E-mini S&P500 futures, where they argue that the substantial depth found at best prices makes it plausible for bets to be proportional to transactions. With modern equity trading strategies, transactions are far smaller than bets (see Hendershott, Jones, and Menkveld (2011) and O'Hara (2015), among others) for typical stocks, and transaction size and frequency varies systematically with trading activity and across trading venues. Moreover, reporting systems differ across trading venues (Upson, Johnson, and McInish (2015)). This means that transaction sizes reported on the consolidated tape convey different information depending on the venue in which the transaction was executed.

higher trading activity.[4] We therefore define trade-time volatility as the volatility over such trade sequences.

We first calculate mean trade-time volatilities of each stock in each month by trading condition (both time-of-day and trading activity level) and then normalize these measures by the corresponding mean trade-time volatility of the median stock. This normalization controls for common variation in the level of trade-time volatility as well as year-month fixed effects. We establish that bet invariance implies that the cross-sectional distributions of normalized trade-time volatilities obtained at different trading conditions should be identical. This means that one can test bet invariance at a given time-of-day window by regressing the cross-section of normalized trade-time volatilities obtained at one trading condition on those obtained at any other trading condition. Invariance predicts that the slope must equal one.

We separately show that the prediction of bet invariance from KO is also implied by a notion that we term "no speculative arbitrage." This analysis considers speculators who maximize expected profits from speculation when they identify a profit opportunity and trades have linear price impacts. We derive the implications of the equilibrium condition that speculators must be indifferent between profit opportunities across different assets and trading conditions. That is, the *expected* profit from investigating one asset rather than another must be the same. Expected profit opportunities are naturally proportional to volatilities. Adjusting for the arrival rates of such opportunities and imposing a no-speculative-arbitrage condition, cross-sections of volatility must be identical across different trading condition environments. When testing at a given time in the trading day, the identifying assumptions made are only solely that the linear price impact of trade of stock j in trading condition x at time-of-day d, λ_{jxd}, and the arrival rate of profit opportunities, β_{jxd}, both separably decompose as $\lambda_{jxd} = \lambda_j \lambda_{xd}$ and $\beta_{jxd} = \beta_j \beta_{xd}$.[5]

An initial regression analysis provides only modest support for either theoretical formulation. Specifically, the bulk of slope coefficient estimates are economically and statistically less than one:

[4]This notion of trading activity differs from that in Kyle and Obizhaeva (2016) in that ours refers to the speed at which a fixed-dollar value is traded as opposed to the arrival time of (unobservable) bets.

[5]When we test no speculative arbitrage across times of day, we must impose the more demanding identification assumptions that $\lambda_{jxd} = \lambda_j \lambda_x \lambda_d$ and $\beta_{jxd} = \beta_j \beta_x \beta_d$.

adjusting type one error for multiple testing, the bet invariance prediction is rejected for over 80% of the 180 possible activity/time-of-day/portfolio combinations. We find a negative association between the magnitudes of departures from predicted slope coefficients of one and goodness of fit. This indicates that departures from bet invariance are likely driven by unexpected variation in volatility, which biases coefficients downward. These departures are the greatest when the independent variable in our tests represents a higher activity level than that of the dependent variable.

The no-speculative-arbitrage condition is, from the speculator's perspective, is an ex-ante concept. That is, speculators allocate resources to acquire information about fundamental values to compete away ex-ante abnormal profits based on their expected profit opportunities (expected volatilities). Our initial tests, in contrast, take realized volatilities as inputs. These realized volatilities reflect both expected profit opportunities and unanticipated better or worse realizations of profits that serve to add noise to the estimates of expected relative profit opportunities. This leads us to decompose the cross-sectional variation in trade-time volatilities entering the right-hand-side of our bet invariance regressions into expected vs. unexpected variation. For each stock in each month, we use the average of its normalized volatilities over the preceding 24 months to measure expected volatilities, and measure unexpected volatility as the difference between the actual and expected quantities.

We find strong support for the predictions of bet invariance when we use expected normalized volatilities as independent variables in our invariance regression tests. With expected volatilities as RHS variables, we reject predictions of bet invariance less than one-third of the time.[6] The magnitudes of departures are barely related to the goodness of fit. This indicates that we largely eliminated the idiosyncratic variation in volatility that drove the departures from invariance found using realized spot volatilities. These results complement findings in KO. They find visual support for invariance, but their formal statistical tests reject invariance. Such rejections may just reflect violations of the stronger auxiliary assumptions imposed by their test construction rather than violations of invariance.

[6]Moreover, the magnitudes of any departures from bet invariance are greatly reduced.

Finally, we extend our analysis to test whether bet invariance holds across time-of-day windows. To do this, we require stronger identifying restrictions requiring decompositions of λ and β by stock, activity level *and* time of day. When using normalized trade-time volatilities from different time-of-day windows as test inputs, estimates are more dispersed and show greater departures from bet invariance. We find large departures from invariance when comparing early and late times in the trading day. In particular, for smaller stocks, coefficient estimates substantially exceed one when the normalized expected volatility on the right-hand side is from late in the trading day and the dependent variable is from early in the trading day, while estimates are far less than one for the opposite scenario. This suggests that attenuation bias is unlikely to drive the documented departures. In sum, we find evidence that speculators equate expected profit opportunities at each moment of the trading day, but not necessarily across different times in the trading day.

The remainder of the paper is organized as follows. Section 2 develops our novel tests of invariance, motivates our use of trade-time volatility, and derives the implications of no speculative arbitrage. Section 3 develops our measure of trade-time volatility, and presents the data. Section 4 presents our findings regarding speculative trading and the cross-section of volatility. Section 5 concludes.

2 Invariance of Risk Transfers and Trade-time Volatility

In this section, we develop a parsimonious test of bet invariance, a fundamental hypothesis underlying MMI. Importantly, our test does not require observability of the exact positions, or "bets", that investors wish to take. We first briefly describe implications of the theory, highlighting relevant assumptions to our analysis. We then reformulate bet invariance using trade-time return volatilities, where volatility is measured with respect to a fixed dollar amount traded of a stock. In doing so, we define the notion of normalized trade-time volatility as the ratio of a stock's trade-time volatility to that of a benchmark stock. Our reformulation of bet invariance shows that this normalized trade-time volatility should have an invariant distribution. We then develop a test, a regression of normalized trade-time volatility at one trading condition on those at other trading conditions, to evaluate bet invariance. We conclude the section by presenting a distinct theoretical framework of speculative trading where speculators maximize expected profits by choosing bet sizes. We show that an equilibrium condition that leaves such speculators indifferent between pursuing profit opportunities across different stocks and trading conditions leads to implications about the cross-section of volatility that nest those implied by bet invariance.

2.1 Invariance of Bets

To formalize MMI, KO introduce the concepts of bet size and bet arrival rates. As described in Andersen et al. (2016), an unsigned bet, denoted $\tilde{Q}_{jt}$, is a random variable that captures a trading decision reflecting a desired change in an investor's position in security j at time t, in shares. Multiplying by the asset price P_{jt}, yields the dollar value, $P_{jt}\tilde{Q}_{jt}$, of this bet. The variable speed at which bets are placed in the market is translated into a bet arrival rate per unit time, denoted N_{jt}. When the calendar-time bet volatility is σ_{jt}^B (in fraction of value per unit time) and the bet arrival rate per calendar-time interval is N_{jt}, the return volatility per bet (per unit of business time) is $\frac{\sigma_{jt}^B}{\sqrt{N_{jt}}}$, presuming that price movements associated with distinct bets are independently distributed. Multiplying the return volatility per bet by $P_{jt}\tilde{Q}_{jt}$ yields the dollar risk transfer associated with implementing a bet,

$$\tilde{I}_{jt} \equiv P_{jt}\tilde{Q}_{jt}\frac{\sigma_{jt}^B}{\sqrt{N_{jt}}}. \tag{1}$$

This object represents the dollar gain or loss associated with implementing the bet $\tilde{Q}_{jt}$. Invariance of bets states that $\tilde{I}_{jt}$ has an invariant distribution across stocks and over time, where the source of random variation in $\tilde{I}_{jt}$ is the variation in bet size $\tilde{Q}_{jt}$.

Interpreting bet volatilities as opportunities to make trading profits, the theory requires that a speculator be indifferent between the potential expected profit opportunities about which she might acquire information. KO argue that were invariance not to hold between two assets, then a trader could more profitably acquire information about the asset that faced larger dollar risk transfers. Essentially, without significant microstructure frictions, traders should compete away abnormal profit opportunities, leaving dollar risk transfers invariant.

Bet size and bet volatility are intrinsically difficult to measure, creating challenges for tests of invariance. Existing empirical analyses sidestep observability concerns in one of two ways: they either (a) use proprietary data for which the econometrician is privy to the exact bets taken (as Kyle and Obizhaeva (2016) assert holds for their data); or (b) assert that within the particular market studied, transactions are proportional to underlying bets (as Andersen et al. (2016) argue holds in the S&P500 E-mini futures market). Our test for invariance of bets addresses this empirical problem. The key feature of our test is that it relies only on the distributional implications of invariance and not the implicit observability of the primitive bets via observation of transactions. Our method translates the invariance of dollar risk transfers of unobservable bets into the invariance of trade-time total return volatility, an *intrinsically observable* object that we define below. Our test uses this observable volatility, as opposed to bet volatility, as an input. Because our approach only relies on the distributional implications of invariance, it enables tests of invariance in a wider variety of markets, in particular those where the bets themselves are unobservable.

2.2 From Bet Volatility to Total Volatility

We begin by distinguishing unobservable bet volume and bet volatility from *observable* "total" volume and "total" volatility, and providing formal links between these notions. Both trading volume and return volatility are affected by sources other than bet implementation. Fundamentally, trading volume is the sum of bet volume and other sources of (e.g., intermediary) volume. Similarly, total return volatility represents aggregate price movements driven by both bet implementation (bet volatility) and other forces, including public information arrival and market frictions.

In addition to requiring the observability of underlying bets, existing tests of invariance such as those in KO make the identifying assumption that both the proportions of bet volume to total volume and bet volatility to total volatility are constant over time and across stocks. Importantly, our test of invariance relaxes these assumptions. Specifically, our method allows for the proportions of bet-to-total volume and bet-to-total volatility to vary with bet size quantile. We require only that this variation maintain a one-to-one mapping from bet volume to total volume and from bet volatility to total volatility. We exploit these one-to-one mappings to derive the implications of bet invariance for the cross-sections of total volatility.

As noted above, existing tests rely on (implicit) observations of bets. Our method translates the invariance of dollar risk transfers of unobservable bets into the invariance of observable trade-time total return volatilities, which we define below. Bet invariance implies that for a fixed stock-specific bet size quantile $\tilde{Q}_{jt}^{(X)}$, the respective invariant $\tilde{I}_{jt}^{(X)}$ (adding bet quantile superscripts to equation (1), above) is constant in the cross-section of stocks and over time.

We relax this structure in two ways. First, we recognize that the arrival rate of information or the amount of information for a stock may vary according to public information characteristics. One can accommodate this formally by imposing a less demanding notion of bet invariance, by only imposing bet invariance conditional on public information C, where C might, for example, capture time of day. This weakens the restrictions implied by bet invariance by only requiring that no speculative arbitrage hold given conditioning information. For example, this allows for the possibility that speculation opportunities may be better in some public information environments than in others. This weaker notion only demands that a speculator be indifferent among the *available* opportunities. This generalization also facilitates controlling for the feature that bet arrival rates likely vary with public information. For example, near open, there may be higher rates of information arrival due to overnight release of information.

Thus, for stocks j and k, at any conditional distribution quantile $x, c \in (0,1)$ (of bet size $\tilde{Q}_{jt}^{(X=x|C=c)}$ and $\tilde{Q}_{kt}^{(X=x|C=c)}$),

$$P_{jt}\tilde{Q}_{jt}^{(X=x|C=c)}\left(\frac{\sigma_{jt}^{B}}{\sqrt{N_{jt}(c)}}\right) \equiv \tilde{I}_{jt}^{(X=x|C=c)} = \tilde{I}_{kt}^{(X=x|C=c)} \equiv P_{kt}\tilde{Q}_{kt}^{(X=x|C=c)}\left(\frac{\sigma_{kt}^{B}}{\sqrt{N_{kt}(c)}}\right). \tag{2}$$

We continue our reformulation of bet invariance conditional on C, omitting the formal dependence on C where it does not cause confusion.

Second, we relax bet invariance to account for the fact that, in practice, the link between bet volume and total volume and that between bet volatility and total volatility may vary systematically. For example, the fraction of total trading volume that is not bet related may be systematically higher or lower around earnings announcements (or certain days of the week, etc.) than at other times, and market frictions may be larger or smaller then, as well.[7] We control for this by allowing

[7] "Parasitic" high-frequency trading is another mechanism that may increase this fraction: high-frequency traders

systematic variation in these market primitives to vary as a function of the bet size quantile x. Importantly, our formulation is agnostic about the exact nature of any such systematic variation. We next formally demonstrate the links between bet volume and total volume, and bet volatility and total volatility, requiring only that bet volume and volatility translate monotonically to total volume and total volatility, respectively.

We define ζ^x to be the ratio of total trading volume to bet volume at bet size quantile x. We capture the ratio of total volatility σ_{jt} to bet volatility σ_{jt}^B by $\psi^x \geq 1$, recognizing that total volatility exceeds bet volatility by $(\psi^x - 1) \times 100$ percent due to news arrival, market frictions, etc. For example, price movements due to the arrival of public macro news induce volatility that speculators may not necessarily be able to bet on.

Our primary identification assumption is that while we allow for variation in these parameters as a function of the bet size quantile x, we assume that $\eta^x = \zeta^x \psi^x$ is a common factor across all stocks. That is, while η^x can vary deterministically in very general ways with bet size quantile x, it does not vary systematically across stocks. With this identification assumption in hand, we multiply equation (2) by η^x. Multiplying the two sides of equation (2) by η^x transforms the bet dollar risk transfer into a function of total dollar volume and total volatility. Hence, we reformulate invariance of bets from MMI into quantitative properties based on observable total volume and volatility. That is,

$$P_{jt}\tilde{Q}_{jt}^{(X=x)}\left(\frac{\sigma_{jt}^B}{\sqrt{N_{jt}}}\right)\eta^x \equiv \tilde{I}_{jt}^{(X=x)}\eta^x = \tilde{I}_{kt}^{(X=x)}\eta^x \equiv P_{kt}\tilde{Q}_{kt}^{(X=x)}\left(\frac{\sigma_{kt}^B}{\sqrt{N_{kt}}}\right)\eta^x, \tag{3}$$

and

$$P_{jt}\tilde{Q}_{jt}^{(X=x)}\zeta^x\left(\frac{\psi^x\sigma_{jt}^B}{\sqrt{N_{jt}}}\right) \equiv \tilde{I}_{jt}^{(X=x)}\eta^x = \tilde{I}_{kt}^{(X=x)}\eta^x \equiv P_{kt}\tilde{Q}_{kt}^{(X=x)}\zeta^x\left(\frac{\psi^x\sigma_{kt}^B}{\sqrt{N_{kt}}}\right). \tag{4}$$

With bet invariance now stated in total volume and volatility terms, we reformulate bet invariance using a fixed dollar amount traded of a stock, eliminating the need to observe individual bets. We consider a constant dollar value S_{qt} for any stock j in a portfolio q of stocks at time t. Because $\tilde{Q}_{jt}$ is stochastic, and bet volume does not represent total volume, we define a random variable $\tilde{A}_{jt}$

who identify institutional order flow remove liquidity in the same direction to profit from providing it at a later time, adding to the amount of non-bet volume and volatility (CITE CITE CITE).

such that

$$S_{qt} = \tilde{A}_{jt} P_{jt} \tilde{Q}_{jt} \zeta. \tag{5}$$

$\tilde{A}_{jt}$ exactly offsets variations in total dollar volume, while ζ scales bet volume, $P_{jt}\tilde{Q}_{jt}$, to total volume. That is, $\tilde{A}_{jt} \equiv \frac{S_{qt}}{P_{jt}\tilde{Q}_{jt}\zeta}$ is the random variable that transforms the constant S_{qt} to the realized dollar bet size, given ζ. For a fixed bet arrival rate N_{jt} and calendar-time bet volatility $\bar{\sigma}_{jt}$, one can interpret $\tilde{A}_{jt}$ as the stochastic time required for dollar volume S_{qt} to trade in the market: $\tilde{A}_{jt}$ is smaller (the required time is shorter) when trading activity is higher. We only require that ζ vary deterministically with bet size quantile in such a way (monotonically suffices) as to preserve a one-to-one mapping between the probability distribution of $\tilde{Q}_{jt}$ and the probability distribution of trading activity (and hence a one-to-one inverse mapping with the probability distribution of $\tilde{A}_{jt}$). As a result, the x^{th} percentile of dollar bet size corresponds to the $1 - x^{th}$ percentile of $\tilde{A}_{jt}$.

From MMI, implementing dollar bet amount $P_{jt}\tilde{Q}_{jt}$ of stock j at time t has a volatility per bet $\frac{\sigma^B_{jt}}{\sqrt{N_{jt}}}$. Equation (4) presents invariance of bets in a form based on total volume and total volatility, implying that the (generally unobservable) business-time *total* volatility associated with trading the rescaled bet value $P_{jt}\tilde{Q}_{jt}\zeta$ is $\frac{\psi\sigma^B_{jt}}{\sqrt{N_{jt}}}$. This leads us to define the observable trade-time total volatility associated with realization of $S_{qt} = \tilde{A}_{jt}P_{jt}\tilde{Q}_{jt}\zeta$ at a given quantile $\tilde{A}^{(X=1-x)}_{jt}$ as

$$\tilde{\sigma}(S_{qt})^{(X=x)}_{jt} \equiv \tilde{A}^{(X=1-x)}_{jt} \left(\frac{\psi^x \sigma^B_{jt}}{\sqrt{N_{jt}}} \right). \tag{6}$$

$\tilde{\sigma}(S_{qt})^{(X=x)}_{jt}$ represents the total volatility associated with realization of S_{qt} dollars worth of stock j's transactions at bet size quantile $\tilde{Q}^{(X=x)}_{jt}$.

2.3 Invariance of Bets and the Cross-section of Total Volatility

We next derive the implications of invariance of bets for the cross-section of trade-time total volatility. The key step is to translate the distributional properties of risk transfers of unobservable bets to distributional properties of observable trade-time volatilities (associated with fixed-dollar values S_{qt}). Equation (5) at the bet size percentile $1 - x$ yields $P_{jt}\tilde{Q}^{(X=1-x)}_{jt}\zeta^{(1-x)} = \frac{S_{qt}}{\tilde{A}^{(X=x)}_{jt}}$. Thus, we can rewrite the invariance condition, equation (4), as:

$$S_{qt}\left(\frac{\psi^{1-x}\sigma_{jt}^{B}}{\tilde{A}_{jt}^{(X=x)}\sqrt{N_{jt}}}\right) \equiv \tilde{I}_{jt}\eta^{(X=1-x)} = \tilde{I}_{kt}\eta^{(X=1-x)} \equiv S_{qt}\left(\frac{\psi^{1-x}\sigma_{kt}^{B}}{\tilde{A}_{kt}^{(X=x)}\sqrt{N_{kt}}}\right). \tag{7}$$

Letting k represent the stock with median bet volatility, we rearrange equation (7) as:

$$\left(\frac{\tilde{A}_{jt}^{(X=x)}}{\tilde{A}_{med,t}^{(X=x)}}\right) = \frac{\psi^{1-x}\sigma_{jt}^{B}\sqrt{N_{med,t}}}{\psi^{1-x}\sigma_{med,t}^{B}\sqrt{N_{jt}}}. \tag{8}$$

Our final step is to multiply both sides of (8) by the right-hand side, so that the left-hand side becomes the ratio of observable trade-time volatilities whose general form was defined in equation (6):

$$\left(\frac{\tilde{A}_{jt}^{(X=x)}\frac{\psi^{1-x}\sigma_{jt}^{B}}{\sqrt{N_{jt}}}}{\tilde{A}_{med,t}^{(X=x)}\frac{\psi^{1-x}\sigma_{med,t}^{B}}{\sqrt{N_{med,t}}}}\right) = \left(\frac{\left(\sigma_{jt}^{B}\right)^{2}}{\left(\sigma_{med,t}^{B}\right)^{2}}\right)\left(\frac{N_{med,t}}{N_{jt}}\right). \tag{9}$$

The left-hand side of equation (9) is what we term *normalized* trade-time volatility (see equation 6),

$$\tilde{Y}_{jt}^{(X)} \equiv \frac{\tilde{\sigma}(S_{qt})_{jt}^{(X)}}{\tilde{\sigma}(S_{qt})_{med,t}^{(X)}} = \left(\frac{\tilde{A}_{jt}^{(X)}}{\tilde{A}_{med,t}^{(X)}}\right)\left(\frac{\sigma_{jt}^{B}\sqrt{N_{med,t}}}{\sigma_{med,t}^{B}\sqrt{N_{jt}}}\right). \tag{10}$$

It is the ratio of stock j's trade-time total volatility to the trade-time total volatility of portfolio q's median stock. The right-hand side of equation (9) is a constant that does not vary with X.

Equation (10) demonstrates the implication of MMI's bet invariance hypothesis for the cross-sectional distribution of observable trade-time volatilities. Specifically, it implies that the trade-time total volatilities associated with S_{qt} at trading activity quantiles (i.e., inverse of $\tilde{A}_{jt}$) x_1 and x_2 must be identical once one normalizes by the median trade-time volatilities at those quantiles.

We have shown that invariance requires that, for all j and t, $\tilde{Y}_{jt}^{(X)}$ be constant in X. Thus, we have that, for distinct activity quantiles x_1 and x_2,

$$\frac{\tilde{\sigma}(S)_{jt}^{(X=x_1)}}{\tilde{\sigma}(S)_{med,t}^{(X=x_1)}} \equiv \tilde{Y}_{jt}^{(X=x_1)} = \tilde{Y}_{jt}^{(X=x_2)} \equiv \frac{\tilde{\sigma}(S)_{jt}^{(X=x_2)}}{\tilde{\sigma}(S)_{med,t}^{(X=x_2)}}. \tag{11}$$

Dividing yields an empirically-testable implication of bet invariance,

$$\frac{\tilde{Y}_{jt}^{(X=x_1)}}{\tilde{Y}_{jt}^{(X=x_2)}} = 1, \quad \forall\ x_1,\ x_2. \tag{12}$$

We test equation (12) using regression analysis and stock-level observations of $\tilde{Y}_{jt}^{(X)}$ obtained at different trading activity conditions. Invariance imposes restrictions not only stock-by-stock, but also with other public information trading features (conditioning information) such as activity level or time-of-day. We extend our formulation to control for time-of-day, denoted $d \in D$ (more generally indexed above by $c \in C$). Invariance imposes that $\tilde{Y}_{jt}^{(X,D)}$ is constant across X and D.

We operationalize the restrictions imposed by invariance by taking natural logs, defining

$$\tilde{y}(X, D)_{jt} \equiv \ln(\tilde{Y}_{jt}^{(X,D)}). \tag{13}$$

To test the strong form of MMI's bet invariance proposed by Kyle and Obizhaeva (2016), we estimate

$$\tilde{y}(x, d)_{jt} = \alpha_0 + \alpha_1 \tilde{y}(x', d')_{jt} + u_{jt}, \tag{14}$$

where $u_{jt} \sim N(0, \delta)$, x and x' index trading activity levels, d and d' index time-of-day windows, and $(x, d) \neq (x', d')$. This strong form of bet invariance predicts that $\alpha_0 = 0$ and $\alpha_1 = 1$. We also consider a weaker form of invariance that only imposes that $\tilde{Y}_{jt}^{(X,D)}$ be constant in X *conditional* on D. To test this we estimate

$$\tilde{y}(x, d)_{jt} = \alpha_0 + \alpha_1 \tilde{y}(x', d)_{jt} + u_{jt}, \tag{15}$$

where $u_{jt} \sim N(0, \delta)$, x and x' index trading activity levels, d indexes a particular time-of-day window, and $x \neq x'$.[8] We next show that the same relationships obtain when profit maximizing speculators are subject to a no-speculative-arbitrage condition.

2.4 No-specuative-arbitrage and the cross-section of total volatility

In this section, we establish that a general version of equation (11), i.e., the implication of bet invariance for the cross-section total volatility, follows from a no-speculative-arbitrage condition. Our derivation starts out with premises that are complementary but distinct from those underlying MMI. When speculators identify a profit opportunity they choose trade amounts that maximize expected profits from speculation. We then impose a no-speculative-arbitrage condition after adjusting for variations in the arrival rates of such opportunities. No speculative arbitrage and bet invariance

[8] In Appendix 4.6, we present estimation results given $d \neq d'$. We find that placing stronger restrictions on how normalized trade-time volatilities must be related across trading condition environments tend to give rise to stronger violations of bet invariance predictions.

yield the same implications for how cross-sections of volatility obtained at different trading conditions must be related. Our development sketches a simple economic mechanism for bet invariance.

No speculative arbitrage is a restriction that says that a speculator should be indifferent as to which stocks she investigates at a moment in time. The price impact of trade for different stocks should adjust in equilibrium to deliver this indifference. While an individual speculator typically only selects among a subset of stocks to investigate, collectively speculators select from the universe of stocks. This delivers the equal expected profit from investigation condition, i.e. the no-speculative-arbitrage condition.

Consider a stock j in market activity condition x at time d of the trading day. The maintained assumption is that if stock j has publicly-known conditional expected value $E[v_j|(x,d)] \equiv E[v_{jxd}]$, then the expected price impact of a trade of size q_{jxd} at that moment takes the familiar linear form

$$p_{jxd}(q_{jxd}, E[v_{jxd}]) = E[v_{jxd}] + \lambda_{jxd} q_{jxd},$$

where the price impact parameter λ_{jxd} can vary with the stock j, the market conditions x (e.g., because the extent of liquidity provision can vary), and the time of day d (e.g., because information arrival can vary). A realized profit opportunity arises when a speculator acquires information that suggests that the true value of the asset is given by $v_{jxd} \neq E[v_{jxd}]$. Optimization leads the speculator to solve

$$\max_{q_{jxd}} q_{jxd}(v_{jxd} - p_{jxd}(q_{jxd}, E[v_{jxd}])) = \max_{q_{jxd}} q_{jxd}(v_{jxd} - E[v_{jxd}] - \lambda_{jxd} q_{jxd}) \tag{16}$$

Solving the resulting first-order condition yields the standard linear trading outcome:

$$q_{jxd} = \frac{v_{jxd} - E[v_{jxd}]}{2\lambda_{jxd}}.$$

Substituting for q_{jxd} into (16) yields the realized profit opportunity,

$$\frac{(v_{jxd} - E[v_{jxd}])^2}{4\lambda_{jxd}}.$$

From an ex-ante perspective, when deciding which stock to investigate, what matters is the expected profit that the speculator anticipates from investigation. This hinges on the expected profit opportunity when one arrives, and the arrival rate of those opportunities.

The expected profit opportunity for stock j conditioned on market activity condition x at time-of-day d is just

$$\frac{E_j\left[(v_{jxd} - E[v_{jxd}])^2 | (x,d)\right]}{4\lambda_{jxd}} \equiv \frac{E_j\left[(v_{jxd} - E[v_{jxd}])^2\right]}{4\lambda_{jxd}} \equiv \frac{\sigma^2_{jxd}}{4\lambda_{jxd}}. \tag{17}$$

That is, the expected profit opportunity is given by stock j's volatility divided by 4 times the price impact of trade. The arrival rate of profit opportunities for stock j in market activity condition x at time d is given by β_{jxd}, where the arrival of profit opportunities can again depend on the stock, level of market activity and time of day. No speculative arbitrage requires that expected profit opportunities, after adjusting for arrival rates of such opportunities, must be equal across stocks and market conditions. That is, for a stock k in market activity condition x' and time-of-day d'

$$\frac{\beta_{jxd}\sigma^2_{jxd}}{4\lambda_{jxd}} = \frac{\beta_{kx'd'}\sigma^2_{kx'd'}}{4\lambda_{kx'd'}} \tag{18}$$

To empirically test whether no speculative arbitrage holds at a given time-of-day window, we make the identifying assumption that the arrival rates of profit opportunities and price impacts decompose as $\lambda_{jxd} = \lambda_j \lambda_{xd}$ and $\beta_{jxd} = \beta_j \beta_{xd}$. That is, there can be heterogeneity in arrival rates of profit opportunities and price impacts across stocks, the impact of trading activity at a given time of day is similar across stocks. Thus, this formulation allows for the possibility that the impact of activity could vary by time of day.

To empirically test the stronger restriction that no speculative arbitrage holds at different time-of-day windows, we make the stronger identifying assumption that the arrival rates of profit opportunities and price impacts decompose as $\lambda_{jxd} = \lambda_j \lambda_x \lambda_d$ and $\beta_{jxd} = \beta_j \beta_x \beta_d$. That is, the impact of trading activity is similar across stocks at different times of the day.

To derive a testable notion of the no-speculative-arbitrage condition at a given time of day, we first observe that it must hold for all stocks including the median stock M in some portfolio q, at any two distinct market activity levels. That is,

$$\frac{\beta_M \beta_{xd} \sigma^2_{Mxd}}{4\lambda_M \lambda_{xd}} = \frac{\beta_M \beta_{x'd} \sigma^2_{Mx'd}}{4\lambda_M \lambda_{x'd}}. \tag{19}$$

Dividing (18) by (19) yields:

$$\frac{\frac{\beta_j \sigma^2_{jxd}}{\lambda_j}}{\frac{\beta_M \sigma^2_{Mxd}}{\lambda_M}} = \frac{\frac{\beta_k \sigma^2_{kx'd}}{\lambda_k}}{\frac{\beta_M \sigma^2_{Mx'd}}{\lambda_M}} \tag{20}$$

Substituting for $j = k$ and using $x \neq x'$, and cancelling the common terms on both sides yields the implication of no speculative arbitrage:

$$\frac{\sigma^2_{jxd}}{\sigma^2_{Mxd}} = \frac{\sigma^2_{jx'd}}{\sigma^2_{Mx'd}}, \tag{21}$$

a general version of equation (11). Our argument shows that with no speculative arbitrage, the cross-section of volatilities relative to the volatility of a benchmark asset, e.g, the asset with median volatility, should not vary across trading conditions. This result is identical to the implication of bet invariance about trade-time volatilities associated with realizations of a fixed dollar volume; we will empirically test this using equation (15) within a given time-of-day window.

With the stronger maintained identifying assumption of separability across stock, activity and time of day, no speculative arbitrage implies

$$\frac{\sigma^2_{jxd}}{\sigma^2_{Mxd}} = \frac{\sigma^2_{jx'd'}}{\sigma^2_{Mx'd'}} \tag{22}$$

for $(x, d) \neq (x'd')$, allowing for tests using activity levels from different times in the trading day. This is a much more demanding notion, both because the maintained identifying assumption is stronger and because it says that speculators can somehow arbitrage intertemporally.

3 Aggregation method and data

3.1 Trade time and normalized trade-time volatility

Our test of bet invariance takes as inputs estimates of volatility that realize as a fixed dollar value S_{qt} is traded (see equations (5)–(15)). We now develop empirical counterparts to $\tilde{y}(x, d)_{jt}$, the log of normalized trade volatility, for a given stock (i), on a given day (t), at a given activity level (x), at a given time of day (d). For each stock, we must first identify trade sequences of fixed-dollar values. We follow the approach suggested by Barardehi, Bernhardt, and Davies (2019).[9] Each year,

[9] Motivated by features of trading in modern markets, Gouriéroux et al. (1999) and Easley et al. (2012), among others, employ similar techniques that aggregate trade information.

we number transactions in stock j sequentially, using index n_j. For transaction n_j, we use $\tau_j(n_j)$, $Q_j(n_j)$, and $P_j(n_j)$ to denote respectively, (i) its time measured in seconds from the beginning of the year, (ii) its size (in shares), and (iii) the corresponding mid-point of best bid and best ask prices.

A trade sequence consists of consecutive transactions that have an aggregate dollar value of at least S_{qt} for each stock in size group q in month t. Thus, a shorter time duration indicates higher trading activity. The first trade sequence begins with the first trade of the year, and each subsequent trade sequence begins with the first trade following the previous sequence.

Formally, we iteratively solve for the last trade of the k^{th} trade sequence, $k = \{1, 2, 3, \ldots\}$, as:

$$n_j^k = \underset{n^*}{\operatorname{argmin}} \left\{ \sum_{n=n_j^{k-1}+1}^{n^*} P_j^C(n) \times Q_j(n) \;\middle|\; \sum_{n=n_j^{k-1}+1}^{n^*} P_j^C(n) \times Q_j(n) \geq S_{qt} \right\}, \tag{23}$$

where $n_j^0 = 0$ and the value of aggregate trades is measured using the previous day's closing price, $P_j^C(n_j)$.[10] Using the previous day's closing price to calculate dollar volumes prevents contemporaneous price movements from affecting identification of trade sequences.[11] We construct trade sequences that span two trading days, but *exclude* them from the analysis. Calculating overnight trade sequences and then excluding them, (1) delivers a random starting point for the first trade sequence of a given day, precluding any systematic bias; (2) circumvents issues associated with overnight price adjustments or information arrival;[12] and (3) avoids combining trading activity levels from near close with those just after open, which typically differ.[13] Figure 1 illustrates a typical pattern.

For stock j, the time duration of the k^{th} trade sequence (i.e., the inverse of trading activity) is,

$$dur_k^j = \tau_j(n_j^k) - \tau_j(n_j^{k-1} + 1). \tag{24}$$

To compute the trade-time volatility in basis points of this trade sequence, we follow the approach suggested by Andersen, Bollerslev, and Diebold (2010) to measure spot volatility and use the square root of the sum of trade-by-trade squared *mid-quote* price returns, where the quoted prices are the

[10] The last quoted bid-ask midpoint is used when the closing price is not available.

[11] Using the previous day's closing price avoids introducing biases driven by contemporaneous price movements. For example, with rapidly increasing prices, using that day's prices can give rise to non-trivially growing dollar volumes, causing a downward bias in the time duration of the corresponding trade sequence.

[12] For instance, we do not need to adjust $Q_j(\cdot)$ and $P_j(\cdot)$ for stock splits or dividend distributions.

[13] Excluding overnight trade sequences drops observations that mix realized volatility from very different trading conditions at close and open of successive trading days. This gives invariance a better chance to succeed in the data.

Figure 1: **Illustration of trade sequence construction.** Panel A presents the dollar volume for a representative mid-sized stock, Acuity Brands Inc. (AYI) on September 22, 2011 over a two hour period. Panel B shows how we aggregate trades into a trade sequence until cumulative dollar volume S_{qt} is reached. As dollar volumes vary over time in Panel A, durations of the associated trade sequence vary in Panel B. Panel C shows evolutions of quoted price mid-point associated with transactions.

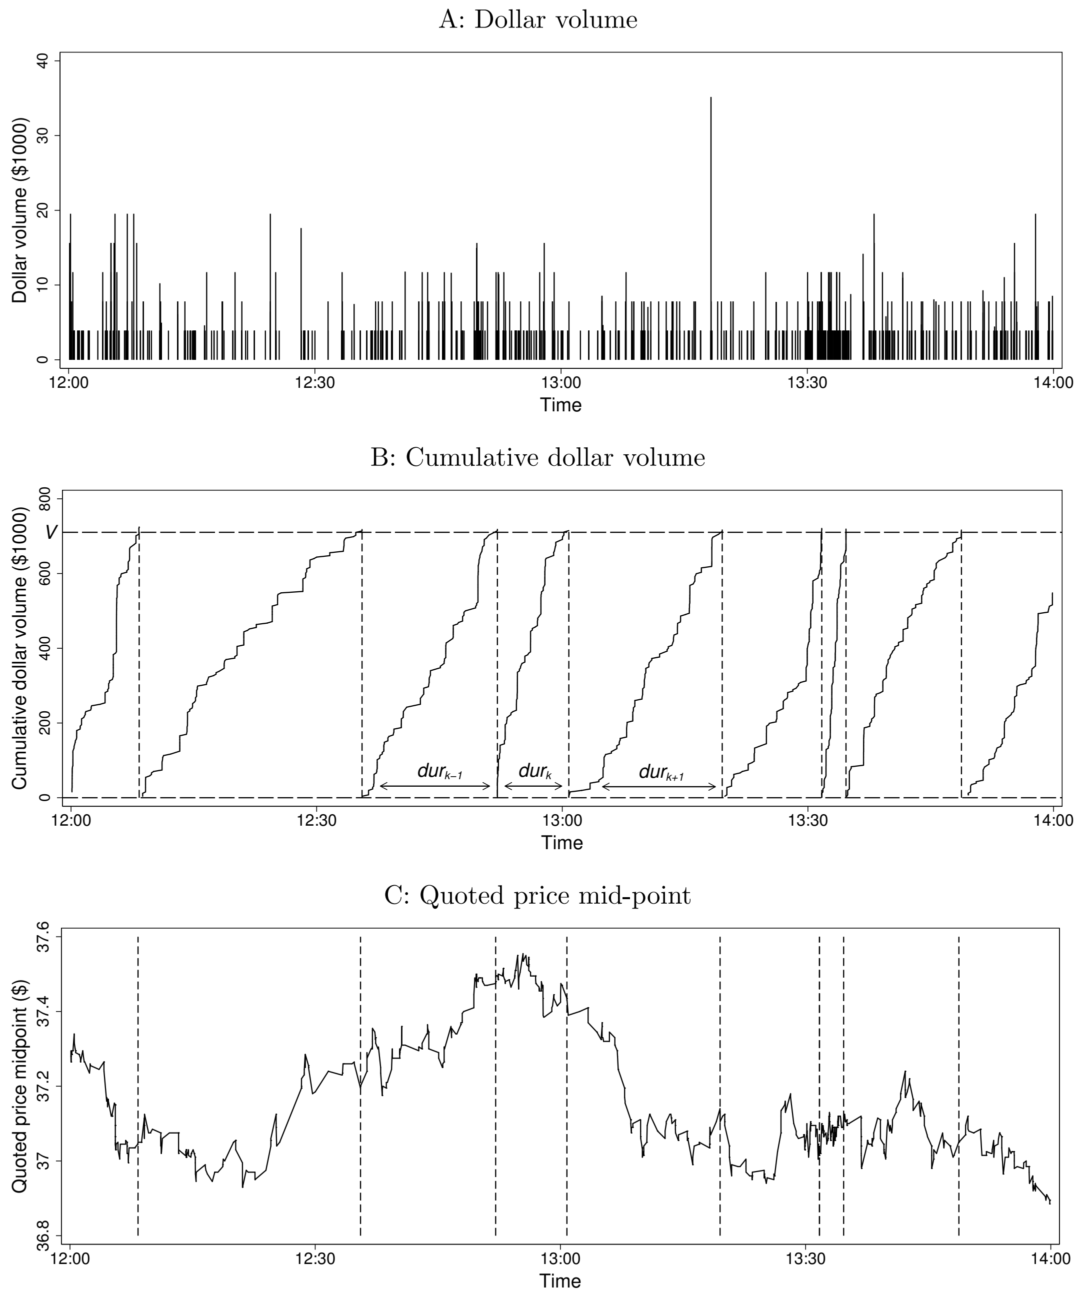

prevailing quotes across all trading venues at the time of each transaction. That is,

$$\sigma_k^j = \sqrt{\sum_{n=n_j^{k-1}+1}^{n_j^k} \left(\frac{P_j(n+1)}{P_j(n)} - 1 \right)^2} \times 10,000. \tag{25}$$

We exclude a trade-by-trade return observation if it exceeds 10% in absolute value. This filter identifies and removes instances of data entry error. This filter matters: a preliminary draft failed to account directly for such data entry errors, which led to a small group of extreme normalized trade-time volatility observations, necessitating a second layer of empirical analysis to account for them.

We next describe how we construct the empirical counterparts of the normalized trade-time volatilities $\tilde{y}(x,d)_{jt}$ (see equation(15)) used in our test of invariance. To begin, we create portfolios of small, mid-sized, and large stocks. Portfolios are formed at the beginning of each year after excluding the largest 30 stocks based on market capitalization at the close of the final trading day of the previous year.[14] The next largest 400 stocks form our portfolio of large stocks, the following 400 stocks form the portfolio of mid-sized stocks, and those with market-cap rankings between 831–1230 form the portfolio of small stocks. Thus, each size portfolio $q \in \{s, m, l\}$ contains 400 stocks. Working with such relatively liquid stocks assures that our findings do not reflect liquid vs. illiquid stock phenomena.

Each month we fix a target dollar value S_{qt} for all stocks in a size portfolio q. The target dollar value is 0.03% of median market capitalization of the stocks in the portfolio at the end of the previous month. This means that we use a total of $3 \times 12 \times 13 = 468$ target dollar values over our 13-year sample period, 2005–2017, or one for each size portfolio by month by year grouping. The average target dollar values for portfolios of large, mid-sized, and small stocks are \$3.098 million, \$0.675 million, and \$0.174 million, respectively. Table 1 shows the evolution of the quartile statistics of time durations of trade sequence that correspond to S_{qt} fixed dollar-values vary over the years in our sample. A typical trade-time observation is 12-30 minutes long, trading off having enough observations to assess intra-daily variations in trading conditions, while still being large enough to avoid the positive autocorrelation associated with dynamic order splitting. In untabulated results, we find that the mean autocorrelation in returns across successive trade sequences is insignificantly different from zero.

[14] We exclude these very large stocks due to the huge variation in their market capitalizations and trading volumes; this heterogeneity means that a fixed targeted dollar value for this group of stocks would be either too small for the largest in the group or too large for smallest stocks in the group. See Ijiri and Simon (1977), Axtell (2001), or Gabaix and Landier (2008) for work on the power law in firm size.

Table 1: **Cross-stock-month averages of stock–specific quartile statistics of trade-time measures associated with 0.03% of median market-capitalization by size portfolio and year.** Trade-time measures, in minutes, of S_{qt} are calculated for each stock. First quartiles, medians, and third quartiles of dur_k^j are calculated for each stock in each month. Cross-stock-month averages of the quartile statistics are reported by stock market-capitalization portfolio and year.

Portfolio	Statistic	Year 2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017
	1^{st} Quartile	16.9	15.9	11.6	8.3	7.5	8.2	8.6	9.5	10.3	10.7	10.7	10.2	11.7
Large	Median	26.9	24.7	18.7	13.3	13.0	14.7	15.2	16.9	18.6	19.0	18.8	18.4	21.4
	3^{rd} Quartile	40.2	36.4	28.0	20.1	20.4	23.4	24.2	26.5	29.4	29.7	29.1	28.7	33.3
	1^{st} Quartile	15.8	14.0	11.2	8.3	9.1	9.5	9.0	9.8	10.2	9.3	9.2	9.4	10.2
Mid-sized	Median	26.5	23.8	18.9	14.3	16.7	18.3	17.7	19.4	20.8	19.0	18.6	19.1	21.9
	3^{rd} Quartile	40.8	36.7	29.8	22.7	27.3	30.6	30.1	32.7	35.6	32.1	31.1	32.2	37.7
	1^{st} Quartile	18.8	16.4	11.9	12.5	14.2	16.4	16.9	21.4	19.6	14.6	15.1	19.7	21.3
Small	Median	32.1	28.6	21.7	23.2	27.2	32.5	34.9	41.9	39.1	30.5	30.9	37.4	41.3
	3^{rd} Quartile	50.4	45.1	35.7	38.7	46.5	55.3	60.5	70.9	66.6	52.7	53.8	61.7	69.4

We assign each trade sequence to one of fifteen time-of-day/activity level categories, first decomposing the trading day into three time-of-day windows, {9:30AM–11:15AM ($d = 1$), 11:15AM–2:15PM ($d = 2$), 2:25PM–4:00PM ($d = 3$)}, and then sorting trade sequences within a time-of-day window into quintiles of trading activity. A trade sequence is assigned to a time-of-day window according to the window containing the mid-point of the trade sequence. For example, a trade sequence beginning at 11:13AM and ending at 11:19AM is assigned to time window $d = 2$. All trade sequences within a month and time of day window are assigned to quintiles of trading activity–trade sequences with shorter time duration feature higher trading activity. These sorts assure an equal number of observations across different activity levels in each time-of-day window of a stock and allow us to control for fixed month effects.

The final step is to calculate the normalized mean trade-time volatility for stock j in month t: $\bar{\sigma}(x,d)_{jt}$ is stock j's mean trade-time volatility, obtained by averaging σ_k^j (equation (25), above) across stock j's trade sequences, in month t at trading activity level x and time-of-day d. The cross-stock median of $\bar{\sigma}(x,d)_{jt}$ in month t for stocks j in portfolio q is denoted $\bar{\sigma}(x,d)^q_{med,t}$. Our empirical analysis focuses on normalized mean trade-time volatility,

$$\bar{y}(x,d)_{jt} \equiv \ln\left(\frac{\bar{\sigma}(x,d)_{jt}}{\bar{\sigma}(x,d)^q_{med,t}}\right) \quad \text{with } j \in q. \tag{26}$$

Normalizing by month and trading condition controls for systematic variation including common time fixed effects. We will show that it controls for the systematic variation in trade-time volatility by time of day and trading activity documented by Barardehi and Bernhardt (2019), and for any common variation in microstructure environments (i.e., in microstructure noise) across trading conditions. Controlling for these factors ensures that we are testing "bet invariance," as delineated in the Market Microstructure Invariance hypothesis in Kyle and Obizhaeva (2016) and as derived in our no-speculative-arbitrage conditions, above, rather than picking up microstructure noise or patterns in volatility associated with trading conditions.[15]

3.2 Data

We focus on the largest U.S.-based NYSE-listed common stocks by market capitalization (CRSP share codes 10 and 11) in our January 1, 2005 to December 31, 2017 sample period. We exclude stocks that do not maintain a minimum daily closing price of $1 over the course of a year. Our sample consists of stocks that feature the necessary identifying information to match with TAQ data. We first use NCUSP from CRSP and CUSIP from TAQ to match data across the two databases; for the stocks without such links, we use TSYMBOL from CRSP and SYMBOL from TAQ to match data. Remaining unmatched stocks are dropped from the sample. As described above, we then drop the 30 largest stocks, and split the next 1200 stocks into subsamples of 400 small, medium, and large stocks. Each month, market-capitalizations are defined as the product of CRSP closing prices and corresponding number of shares outstanding at the end of the previous month.

Transaction quantities and time stamps are obtained from the consolidated trade history in the NYSE TAQ database. We consider all stock trades on all U.S.-based trading venues, including those regulated by FINRA, during regular market hours between 9:30AM and 4:00PM (EST). We construct National Best Bid and National Best Offer prices at the time of each transaction from the consolidated quotes history and NBBO files in the TAQ database, and match each transaction with the corresponding mid-point price.[16]

[15] Our findings are unaffected by the benchmark used for normalization. In untabulated robustness analyses, identical findings obtain when, instead of the stock with month-specific median volatility, we normalize trade-time volatilities with respect to the volatility of the stock that is at the first or the third quartile of a given monthly cross-section.

[16] We use a modified version of the SAS code available on Professor Craig Holden's website

4 Cross-section of trade-time volatility

In this section, we conduct our test of bet invariance (no-speculative-arbitrage) in the cross-section of trade-time volatilities. We first establish that normalization by the median stock in each sample eliminates all economically meaningful time-of-day patterns in realized normalized trade-time volatilities. We first regress these realized volatilities on realized volatilities across trading activity levels, and find substantial departures from the central prediction of bet invariance. Analysis of these departures reveals a systematic relationship: departures grow larger when comparing less active to more active markets, and the corresponding goodness of fit in these regressions weakens. Because, the notion of no speculative arbitrage is an ex-ante concept—a speculator's decision about which asset to investigate depends on the expected profit they anticipate from investigation—realized volatilities, as independent variables, possess noise that biases coefficients downward. To correct for this, we next decompose volatilities on the right-hand sides of our invariance regressions into expected and unexpected volatilities. This corrected analysis provides strong support for bet invariance, revealing that the initial departures were largely driven by unexpected cross-sectional variation in realized volatilites.

4.1 Time-of-day effects in the cross-section of trade-time volatility

We first show that normalizing trade-time volatilities with respect to a benchmark stock—the stock with median volatility—controls for level effects of trading activity and time-of-day on trade-time volatility.[17]. This normalization avoids conflating variation driven by level effects with the relative cross-sectional differences in profit opportunities associated with bet invariance.

Figure 2 contrasts the cross-sections of log *actual* trade-time volatilities with those of *normalized* trade-time volatilities. The top panel plots quartile statistics averaged over months of log actual trade-time volatilities for different time-of-day windows and trading activity levels. The figure shows that actual trade-time volatility drops over the trading day and also with trading activity. In contrast, the bottom panel reveals that normalized trade-time volatilities, $\bar{y}(x,d)_{jt}$, do not mean-

(https://kelley.iu.edu/cholden/) to construct best quoted prices and match midpoints with transactions.

[17]See Barardehi and Bernhardt (2019) for a discussion of these patterns.

Figure 2: **Cross-section of trade-time volatility by size portfolio.** For each market-cap category in each month medians, 25^{th} percentiles, and 75^{th} percentiles of trade-time volatility ($\ln(\bar{\sigma}(x,d)_{jt})$ and the normalized trade-time volatility ($\bar{y}(x,d)_{jt}$) are calculated. Averages of the three statistics are obtained by time-of-day window d and trading activity level x across the 48 months of data. Market-capitalization portfolios are formed at the beginning of the year.

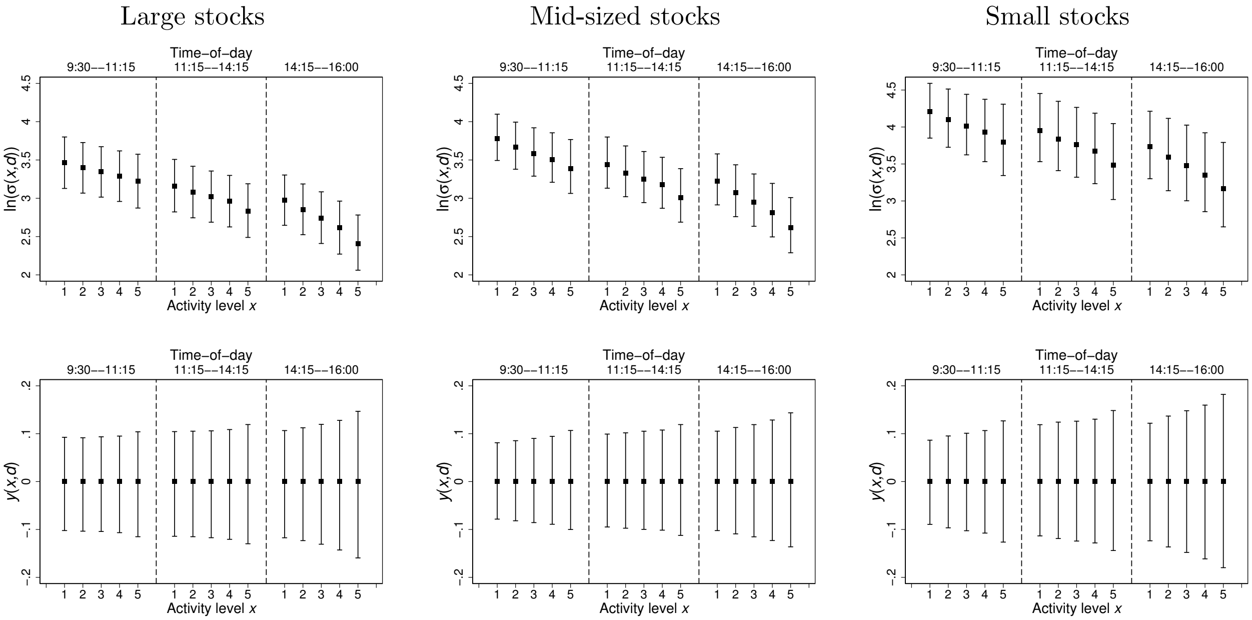

ingfully vary with trading activity or time of day—normalization removes almost all of the level effects associated with different trading conditions. Indeed, the cross-stock medians across trading activity levels and time of day are virtually identical, and the other quartile statistics display only modest variation as activity rises.

4.2 Preliminary test of bet invariance

We now test the less demanding form of bet invariance as posited by KO and as formulated above in our no-speculative-arbitrage conditions. To do so, we estimate the linear association between normalized mean trade-time volatility at a given time-of-day and activity level, from equation (26), and that in all other activity levels in the same time-of-day window. That is, we estimate

$$\bar{y}(x,d)_{jt} = \alpha_0 + \alpha_1 \bar{y}(x',d)_{jt} + u_{jt} \quad \text{with } u_{jt} \sim N(0,\delta), \qquad \text{for} \quad x, x' \in \{1,2,3,4,5\} \quad \text{and} \quad d \in \{1,2,3\}. \tag{27}$$

With five activity levels, indexed by $x \neq x'$, and three time-of-day windows, indexed by d, there are a total of $(5 \times 4) \times 3 = 60$ ordered pairs of dependent and independent variables, reflecting

$\frac{5\times4}{2}\times 3 = 30$ distinct unordered pairs of trading activity levels, accounting for different time-of-day, for equation (27). For each stock size category q, we estimate equation (27) for each of these 60 pairs using a Fama-MacBeth approach.[18] We then investigate estimated α_1 coefficient magnitudes and their statistical significance, and formally test the prediction of bet invariance, $H_0 : \alpha_1 = 1$, each time we estimate equation (27).

Figure 3 shows that, at a first glance, the data do not provide overwhelming support for the prediction of bet invariance. The top row presents histograms for $\hat{\alpha}_1$. The histograms indicate that most point estimates fall below one: point estimates range between 0.68 and 1.04, and they tend to fall further below for smaller and mid-sized stocks.[19] The middle row shows the histograms of the corresponding t-statistics that test the null of $H_0 : \alpha_1 = 1$. The overwhelming majority of estimates are statistically different from one (most estimates are less than one, and over 80% of the t-statistics exceed 4 in absolute value, with over 95% exceeding 2). Of course, when we formally test the hypothesis $H_0 : \alpha_1 = 1$, we reject the null most of the time—but these differences are not quantitatively large. Invariance also predicts that $\alpha_0 = 0$. Consistent with this, we find that the economic magnitudes of α_0 estimates are negligible in magnitude, varying between -0.027 and 0.025, even though nearly half of these estimates are statistically different from zero and only 20% of estimates have t-statistics that exceed 4 in absolute value. The bottom row in Figure 3 shows that for all size portfolios and in all years, equation (27) fits the data well, with most R^2 magnitudes exceeding 80%. Our initial findings do not support the theory of bet invariance or a no-speculative-arbitrage condition.

We examine all 60 ordered pairs of normalized trade-time volatility that one can choose from the five trading activity levels, within different three time-of-day windows. While there are only 30 distinct unordered such pair sets, theory does not specify the choice of dependent and independent variables in the regression analysis. We next show that this choice matters, as it explains the magnitudes of violations of bet invariance.

[18] We display Newey-West standard errors.

[19] This of course refers to smaller (rank 701–1000) stocks within the largest 1000 stocks.

Figure 3: **Empirical distributions of $\hat{\alpha}_1$s, $t(H_0 : \alpha_1 = 1)$s, and R^2s by stock size category.** For each market-cap category, equation (27) is estimated using a Fama-MacBeth approach given the 60 possible pairs of normalized mean trade-time volatilities at different time-of-day windows and activity levels ($\bar{y}(x, d)$ and $\bar{y}(x', d)$). The top row presents the empirical distributions of $\hat{\alpha}_1$ point estimates. The middle row presents the empirical distributions of corresponding t-statistics for $H_0 : \alpha_1 = 1$, with vertical dashed lines indicating $t = -4$ and $t = 4$. The bottom row presents the empirical distributions of R^2s.

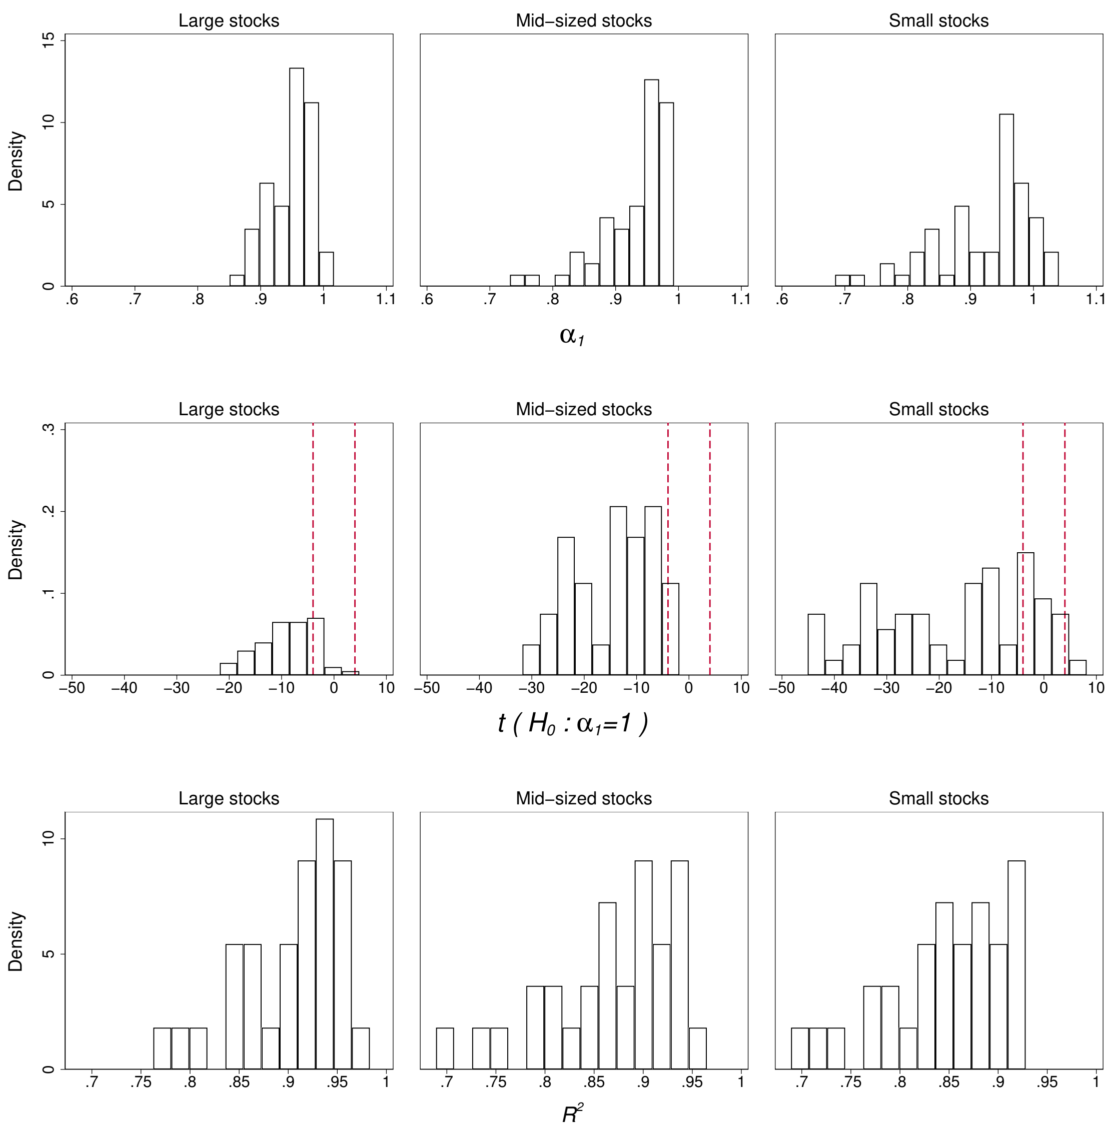

4.3 Trading activity level and the goodness of fit

We next explore the magnitude of departures from bet invariance, i.e., the amount by which α_1 estimates differ from unity (shown in Figure 3). We begin by examining the association between departures from bet invariance and measures of goodness of fit for equation (27). We then relate these associations to the levels of trading activity that underlie pair sets of normalized trade-time volatility entering our tests.

Figure 4: **Co-variation between $\hat{\alpha}_1$ and R^2 by trading activity levels.** For each market-cap category, equation (27) is estimated using a Fama-MacBeth approach given the 60 possible pairs of normalized mean trade-time volatilities at different time-of-day windows and activity levels ($\bar{y}(x,d)$ and $\bar{y}(x',d)$). The figures present the correspondence between $\hat{\alpha}_1$ and R^2 from the fit based on each pair. Shapes are color-coded to gray and black to identify $\bar{y}(x,d)$ and $\bar{y}(x',d)$ pairs with $x > x'$ and $x < x'$, respectively. Trading activity levels $x, x' \in \{2,3,4,5\}$ are reflected by symbols { + , △ , □ , ○ }, respectively.

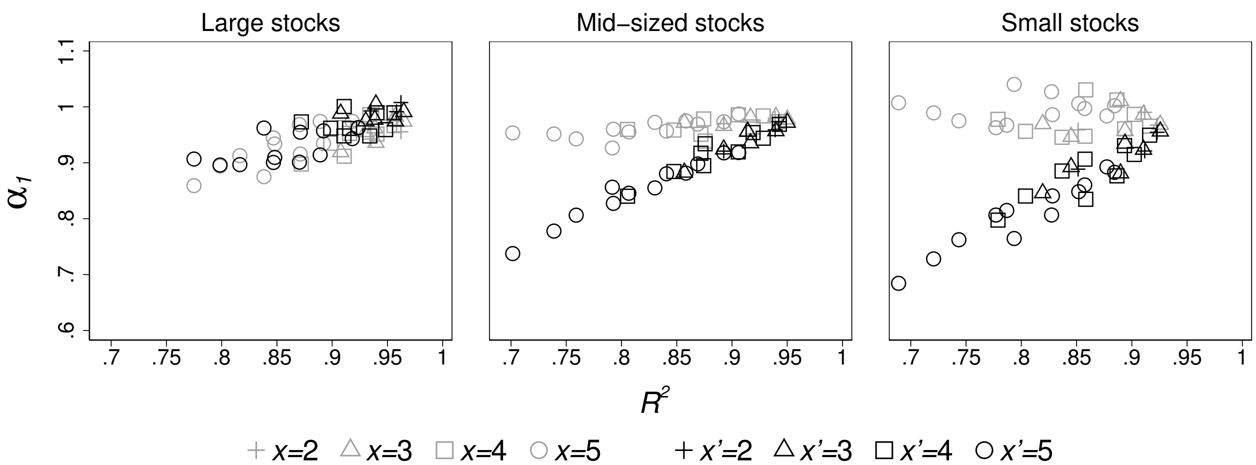

Figure 4 illustrates two important findings. First, it shows that estimates of $\hat{\alpha}_1$ from equation (27) tend to be smaller when R^2 is smaller. As shown by the data points in Figure 4, R^2 is far smaller when higher trading activity levels x or x' underlie trade-time volatilities used as inputs. Second, for portfolios of small and mid-sized stocks, this effect is shown predominantly in one subsample of ordered activity level ordered pairs: when $x > x'$, i.e., when the independent variable corresponds to the less active market condition, $\hat{\alpha}_1$'s vary minimally with R^2 as x rises and are all close to one; in contrast, for $x < x'$, $\hat{\alpha}_1$ estimates fall sharply along with their corresponding R^2's as x' rises, especially when x' corresponds to the most active market conditions. This pattern intensifies going from the portfolio of mid-sized firms to that of small firms. In other words, departures from bet invariance tend to reflect situations where the level of trading activity x' that underlies normalized trade-time volaltilities entering the right-hand-side (RHS) of equation (27) is higher.

These results are consistent with idiosyncratic components driving more of the cross-sectional variation in volatilities in more active markets. For example, such idiosyncratic variation in volatilities can reflect public information news arrival that drives price movements and market activity. Such increased idiosyncratic variation should naturally reduce measures of goodness of fit. That is, when high idiosyncratic variation enters a RHS measure of volatility, it biases estimates of α_1 downward, as they are a component of volatility that is not orthogonal to expected speculative profit opportunities. These patterns are weaker for the portfolio of large stocks, where α_1 coefficients are closer to one, R^2s are higher, and trade-time volatilities obtained at higher trading activities tend to produce smaller $\hat{\alpha}_1$ coefficients regardless of the relationship between x and x'. This result suggests that idiosyncratic variation is relatively greater for smaller, less actively traded and followed, stocks. We next show support for this conjecture by splitting the cross-sectional variation in volatility into systematic and idiosyncratic components.

4.4 Reconsidering expected profit opportunities

Bet invariance, in our setting, delineates a no-speculative-arbitrage condition in the following way: the profit opportunities from speculation in a given asset, relative to those in a benchmark asset, e.g., the median stock, should be constant across trading conditions. With free entry, speculators arbitrage away any differences across trading conditions. This interpretation does not preclude large *realized* differences between profit opportunities across stocks due to unanticipated shocks or public information arrival that may increase or reduce volatility, that result in deviations from predictions of bet invariance. The distinction between expected versus realized profit opportunities is particularly important because of our empirical design. Our monthly estimates of trade-time volatility measure *spot* volatility (see equation (25)); and spot volatility is sensitive to high-frequency, e.g., intraday, price fluctuations. Thus, normalized trade-time volatilities can vary due to unanticipated information- or liquidity-driven price movements that do not enter the ex-ante considerations of speculators, as they are unlikely to be exploitable by speculators.

To obtain a better measure of expected profit opportunities, we now decompose the cross-sectional variation in trade-time volatility at a given trading activity level and time-of-day window into *expected* vs. *unexpected* volatilities. Expected variations in trade-time volatility proxy the differences in expected profit opportunities that speculators ex ante consider when allocating resources to acquire information about fundamental values in the universe of assets. Realized cross-stock variation in normalized trade-time (spot) volatility captures both public information arrival from which speculators cannot profit, and potentially unexpected variation in profit opportunities. This unexpected component is an individual stock's improved or worsened speculation opportunities, relative to a benchmark asset. In sum, the normalized trade-time (spot) volatilities that the econometrician uses measures with noise the true relative expected profit opportunities that enter speculator decision-making and the no-speculative arbitrage condition, especially true in active market conditions. To address this, we now estimate expected profit opportunities over time to reduce the noise in our measure.

To integrate this theoretically, we take expectations *over time* of the two sides of the invariance relation, equation (11):

$$\ln\left(\tilde{Y}_{jt}^{(x,d)}\right) = \ln\left(\tilde{Y}_{jt}^{(x',d')}\right) \quad \Rightarrow \quad E_t\left[\ln\left(\tilde{Y}_{jt}^{(x,d)}\right)\right] = E_t\left[\ln\left(\tilde{Y}_{jt}^{(x',d')}\right)\right]. \tag{28}$$

We implement this empirically by using the mean of the previous 24 months of spot trade-time volatilities as an estimate of $E_t\left[\ln\left(\tilde{Y}_{jt}^{(x,d)}\right)\right]$.

Assuming an additive error structure for the LHS of equation (28), we have

$$\ln\left(\tilde{Y}_{jt}^{(x,d)}\right) + U_{jt}^{(x,d)} = E_t\left[\ln\left(\tilde{Y}_{jt}^{(x',d')}\right)\right]. \tag{29}$$

Thus,

$$\ln\left(\tilde{Y}_{jt}^{(x,d)}\right) = E_t\left[\ln\left(\tilde{Y}_{jt}^{(x',d')}\right)\right] - U_{jt}^{(x,d)}, \tag{30}$$

where $U_{jt}^{(x,d)}$ is a mean-zero random innovation.

We posit that the departures documented in Figure 3 reflect the distinction between expected profit opportunities that speculators pursue ex ante vs. an unpredictable component that does not enter a speculator's decision making.

4.5 Bet invariance of expected speculative opportunities

To empirically capture the two sources of cross-stock variation in profit opportunities to speculation, we decompose equation (27)'s RHS trade-time volatility into expected vs. unexpected components.

We use each stock's average normalized trade-time volatility over the previous 24 months (or the longest time-series available within the 24 months) to proxy for expected volatilities. The difference between the realized and expected volatilities captures unexpected volatility. That is, expected trade-time volatility is

$$\bar{y}(x', d)^E_{jt} = \frac{1}{L_{jt}} \sum_{m=t-25}^{t-1} \bar{y}(x', d)_{jm}, \tag{31}$$

where L_{jt} is the number of $\bar{y}(x', d)_{jm}$ observations between months $t-1 \leq m \leq t-25$. Accordingly, unexpected volatility is given by

$$\bar{y}(x', d)^U_{jt} = \bar{y}(x', d)_{jt} - \bar{y}(x', d)^E_{jt}. \tag{32}$$

We operationalize an empirical test of equation (30) by replacing the RHS variable in equation (27) with $\bar{y}(x', d)^E_{jt}$ and $\bar{y}(x', d)^U_{jt}$, separately, to investigate bet invariance by examining the cross-stock distributions of expected vs. unexpected normalized volatilities, respectively.

We now show that the prediction of bet invariance is largely retrieved when we use expected normalized trade-time volatilities as the independent variable in our bet invariance regression. We establish that the departures from the implications of bet invariance are almost entirely driven by unexpected variation in the cross-section of trade-time volatilities. We estimate

$$\bar{y}(x, d)_{jt} = \alpha^E_0 + \alpha^E_1 \bar{y}(x', d)^E_{jt} + u^E_{jt} \quad \text{with } u^E_{jt} \sim N(0, \delta), \tag{33}$$
$$\text{for} \quad x, x' \in \{1, 2, 3, 4, 5\} \quad \text{and} \quad d \in \{1, 2, 3\}$$

and

$$\bar{y}(x, d)_{jt} = \alpha^U_0 + \alpha^U_1 \bar{y}(x', d)^U_{jt} + u^U_{jt} \quad \text{with } u^U_{jt} \sim N(0, \delta), \tag{34}$$
$$\text{for} \quad x, x' \in \{1, 2, 3, 4, 5\} \quad \text{and} \quad d \in \{1, 2, 3\},$$

maintaining the estimation approach used to fit equation (27).

Figure 5 illustrates support for bet invariance when we substitute expected volatilities for realized volatilities on the right-hand side of the invariance regressions. Compared to Figure 3, the histograms of the slope coefficients, α^E_1, display visibly denser concentrations near 1 for all stock

Figure 5: **Empirical distributions of $\hat{\alpha}_1^E$s and $t(H_0 : \alpha_1^E = 1)$s.** For each market-cap category, equation (33) is estimated using a Fama-MacBeth approach given the 60 possible pairs of normalized mean trade-time volatilities across different time-of-day windows and at different activity levels, ($\bar{y}(x, d)$ and $\bar{y}(x', d)$). The top row presents the empirical distributions of $\hat{\alpha}_1$ point estimates. The bottom row presents the empirical distributions of corresponding t-statistics for $H_0 : \alpha_1 = 1$, with vertical dashed lines indicating $t = -4$ and $t = 4$.

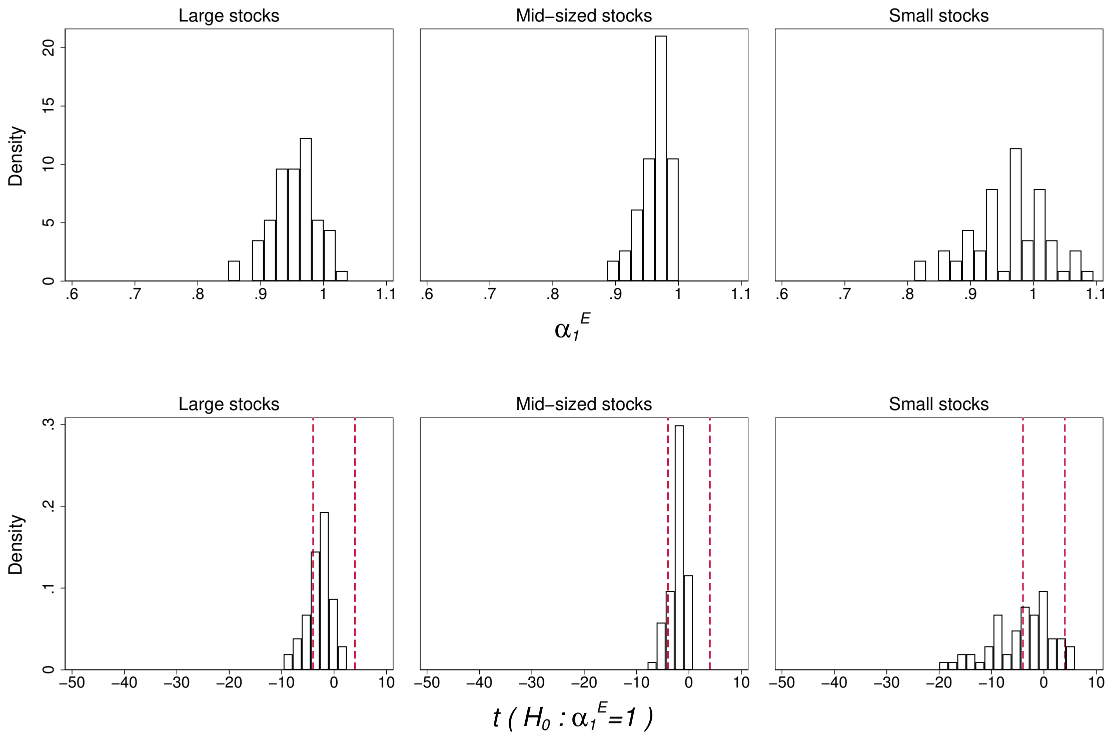

size portfolios. Point estimates now range between 0.81 and 1.09, indicating an upward density shift relative to those shown in Figure 3. Moreover, the t-statistics testing the null $H_0 : \alpha_1^E = 1$ are far more concentrated around zero than those for $H_0 : \alpha_1 = 1$ as in our previous tests. Only 29% of t-statistics exceed 4 in absolute value, with just over 60% exceeding 2. As a result, one cannot reject the prediction of bet invariance in most cases when we use this improved measure of expected profit opportunities.

Moreover, the use of expected volatilities as the right-and side variable in invariance regressions largely eliminates the association between deviations from bet invariance and measures of goodness of fit. Figure 6 plots α_1^E estimates against the corresponding R^2 measures. The patterns reinforce our premise that deviations from the prediction of bet invariance found using normalized spot volatilities reflect that these normalized spot volatilities are noisy measures of expected volatilities. Once much of the measurement error is removed, not only do the slope coefficients shift upward, but the association between departures from invarance and R^2 measures of invariance regressions largely vanish, especially for mid-sized and large stocks.[20]

Figure 6: **Co-variation between $\hat{\alpha}_1^E$ and R^2 by trading activity levels.** For each market-cap category, equation (33) is estimated using a Fama-MacBeth approach given the 60 possible pairs of normalized mean trade-time volatilities at different time-of-day windows and activity levels ($\bar{y}(x,d)$ and $\bar{y}(x',d)$). The figures present the correspondence between $\hat{\alpha}_1^E$ and R^2 from the fit based on each pair. Shapes are color-coded to gray and black to identify $\bar{y}(x,d)$ and $\bar{y}(x',d)^E$ pairs with $x > x'$ and $x < x'$, respectively. Trading activity levels $x, x' \in \{2,3,4,5\}$ are reflected by symbols { + , △ , □ , ○ }, respectively.

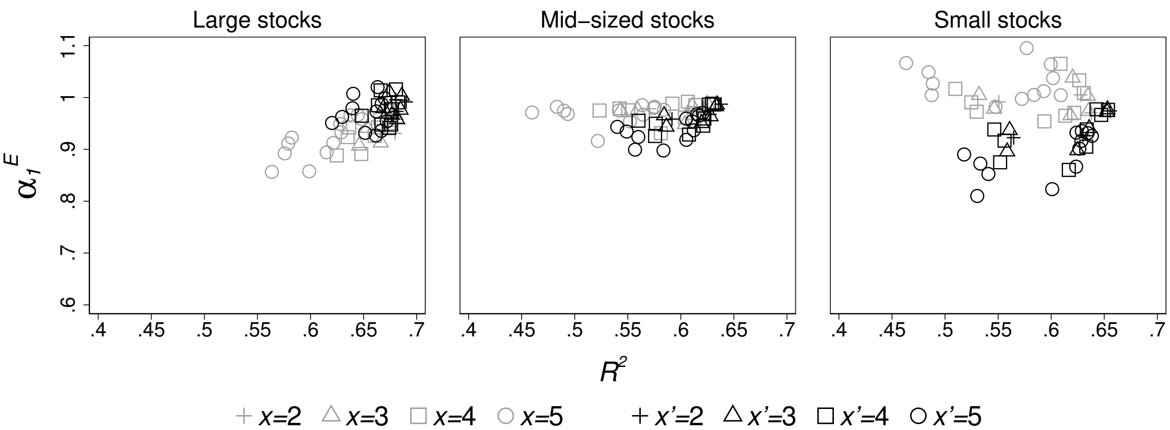

We conclude our analysis by providing more direct evidence that the deviations from bet invariance are driven by realizations of unexpected components of realized normalized volatilities. To do this, we establish the strict stochastic dominance of the distributions of $\hat{\alpha}_1^U$ with respect to those of $\hat{\alpha}_1^E$. Figure 7 demonstrates this graphically: cumulative densities of $\hat{\alpha}_1^U$ substantially exceed those of $\hat{\alpha}_1^E$ for all values of α_1 estimates. Averages of $\hat{\alpha}_1^U$ for portfolios of large, mid-sized, and small firms are 0.73, 0.78, and 0.75 respectively; whereas the counterparts for $\hat{\alpha}_1^E$ are 0.95, 0.96, and 0.96, demonstrating the fundamental effect of unexpected variation in volatility in driving the departures from bet invariance.

[20]The R^2 measures in Figure 6 are 20-30 percentage points less than their counterparts in Figure 4. These differences reflect that equations (33) and (27) share the same dependent variable, but the variation in the independent variable of equation (33) is smaller. In fact, $\bar{y}(x',d)_{jt} = \bar{y}(x',d)_{jt}^E + \bar{y}(x',d)_{jt}^U$, and $\bar{y}(x',d)_{jt}^U$ has a mean of zero. Thus, the numerator of the R^2 statistic for equation (33) is, by construction, smaller than that for equation (27), while the denominators are the same.

Figure 7: **Cumulative densities of $\hat{\alpha}_1^U$ and $\hat{\alpha}_1^E$.** For each market-cap category, equations (33) and (34) are estimated using a Fama-MacBeth approach given the 60 possible pairs of normalized mean trade-time volatilities at different time-of-day windows and activity levels. For each size portfolio, the cumulative density distributions are plotted.

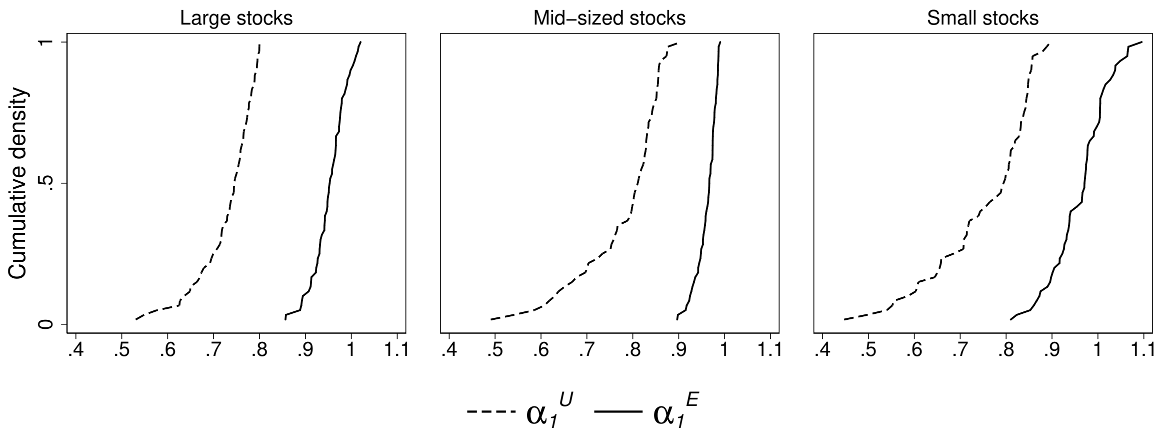

4.6 Bet invariance of expected speculative opportunities across time-of-day

We now extend our analysis to investigate whether there are abnormal profit opportunities across time-of-day. To do this, we must make the stronger identifying assumption that $\lambda_{jxd} = \lambda_j \lambda_x \lambda_d$ and $\beta_{jxd} = \beta_j \beta_x \beta_d$ (see equation (18)). We use normalized volatilities obtained from *different* time-of-day windows, estimating

$$\bar{y}(x,d)_{jt} = \alpha_0^E + \alpha_1^E \bar{y}(x',d')_{jt}^E + u_{jt}^E \quad \text{with } u_{jt}^E \sim N(0,\delta), \tag{35}$$

$$\text{for} \quad x, x' \in \{1,2,3,4,5\} \quad \text{and} \quad d, d' \in \{1,2,3\}.$$

This specification is the same as that in equation (33), save that x and x', indexing trading activity levels, may or may not be equal, and $d \neq d'$ index time of day. With three times of day and five activity levels and the restriction that time of day must differ on the two sides of equation (35), we have a total of $(3 \times 5) \times (2 \times 5) = 150$ regressions. Figure 8 shows that using expected trade-time volatilities as independent variables when testing the stronger form of bet invariance does *not* retrieve the unit coefficient prediction of no speculative arbitrage. To the contrary, α_1^E estimates given different time-of-day display larger dispersion than those obtained within time-of-day.

That predictions of a stronger form of bet invariance are not supported by the data suggests

Figure 8: **Cumulative densities of $\hat{\alpha}_1^E$ for $d = d'$ vs. $d \neq d'$.** For each market-cap category, equations (33) and (35) are estimated using a Fama-MacBeth approach given the, respectively, 60 and 150 possible pairs of normalized mean trade-time volatilities at different time-of-day windows and activity levels. For each size portfolio, the cumulative density functions of $\hat{\alpha}_1^E$ obtained from the two specifications are plotted.

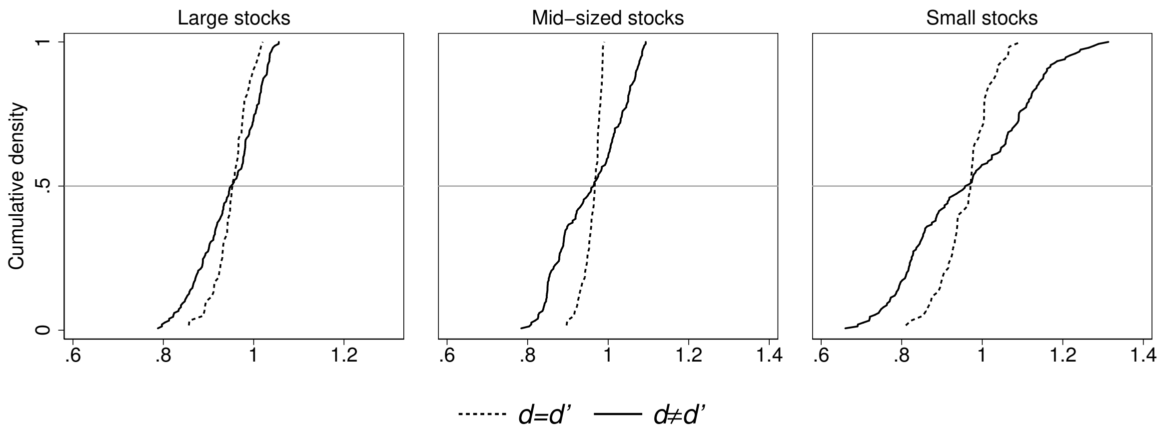

that speculative arbitrage may not be carried out across different times of the trading day. We provide supportive evidence of this by conditioning α_1^E estimates on the distance between the time-of-day windows underlying the respective LHS and RHS normalized volatility measures. That is, we condition estimates according to whether or not both LHS and RHS volatilities in our tests represent early or late trading hours. In this way, we isolate LHS and RHS combinations that represent the largest time-of-day disparities. Figure 9 shows that the greatest departures from $\alpha_1^E = 1$ are present in cases in which the equation (35) LHS is from early in the day and the RHS is from late in the day. Furthermore, the greatest dispersion in estimates is found when including midday volatilities on either side of the regression. This provides further evidence that speculative arbitrage may be more effective within a smaller interval of time.

Analysis in Section 4.5 revealed that expected normalized trade-time volatilities at a given time of the day are roughly equal regardless of the activity level, consistent with the implication of no-speculative arbitrage. That is, at any given time in the trading day, speculators adjust their resource allocation to deliver roughly an invariant distribution of expected returns to speculation across stocks. Also recall that our test of bet invariance jointly tests the identifying restrictions imposed on price impacts (λ) and the arrival rates of speculative profit opportunities (β). As

Figure 9: **Cumulative densities of $\hat{\alpha}_1^E$ when or $d = 1$ and $d' = 3$ vs. $d = 3$ and $d' = 1$.** For each market-cap category, equation (35) is estimated using a Fama-MacBeth approach given the 150 possible pairs of normalized mean trade-time volatilities at different time-of-day windows and activity levels. For each size portfolio, the cumulative density functions of $\hat{\alpha}_1^E$ obtained given $d = 1$ and $d' = 3$ vs. $d = 3$ and $d' = 1$ are plotted.

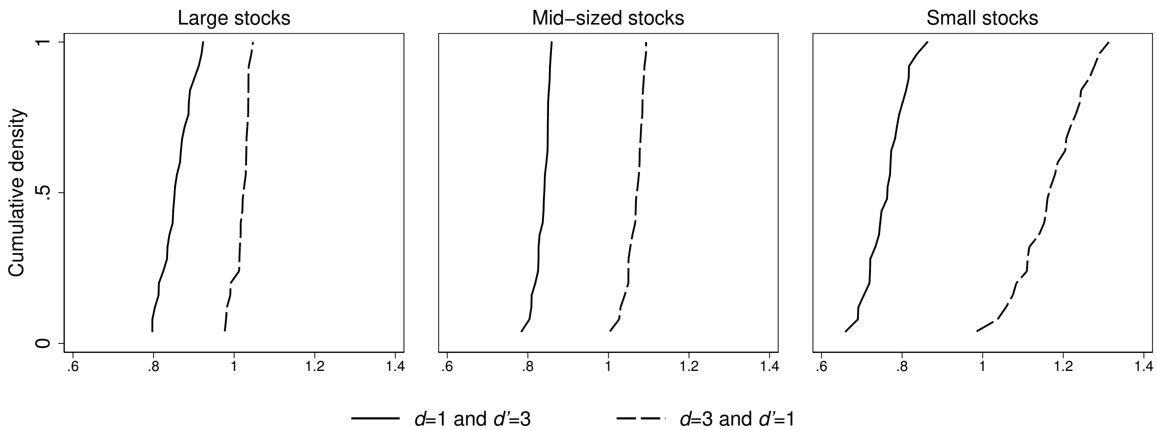

such, finding empirical support for bet invariance (no speculative arbitrage) within time-of-day windows indicates the plausibility of our identifying assumptions that the linear price impact of trade and the arrival rate of profit opportunities decompose as $\lambda_{jxd} = \lambda_j \lambda_{xd}$ and $\beta_{jxd} = \beta_j \beta_{xd}$. To extend bet invariance or our notion of no speculative arbitrage across different times of day requires stronger identifying restrictions, requiring decomposition across stocks, activity levels *and* time of day, $\lambda_{jxd} = \lambda_j \lambda_x \lambda_d$ and $\beta_{jxd} = \beta_j \beta_x \beta_d$.

Figure 9 shows that, when d and d' are from early and late in the day, respectively, coefficient estimates vary widely from the predicted value of 1. For smaller stocks, within our sample of relatively larger stocks, this is also true when d is late, and d' is early. Moreover, it indicates that attenuation bias does not underlie deviations of invariance from one for small stocks: α_1^E estimates substantially exceed one for small stocks when $d = 3$ and $d' = 1$, and any attenuation bias would only reduce these estimates. Because measurement errors are more likely to affect results for smaller stocks, this also suggests that attenuation bias is unlikely to drive estimates that are less than one when $d = 1$ and $d' = 3$. One must still be cautious when interpreting these estimates. Due to the necessary identifying restrictions for testing across times of day, these findings need not reflect violations of invariance. It could instead be that either the price impacts or arrival rates of profit opportunities feature stock-specific variations across different points in a trading day, or both.

5 Conclusion

This paper builds upon and extends existing market microstructure theories to test whether speculators arbitrage away abnormal profit opportunities across assets and trading conditions. Our paper derives testable implications from the Market Microstructure Invariance hypothesis "bet invariance," developed by Kyle and Obizhaeva (2016), for the cross-section of volatility that rely only on publicly-available data. Our test addresses the non-observability of bets by translating implications for risk transfers of bets into implications for the distribution of observable trade-time volatilities. Moreover, our framework is less restrictive than previous tests of invariance in that our formulation allows for heterogeneous arrival of bets and price impacts on stocks, over time and across trading activity levels. We show that bet invariance implies that the slope of a cross-sectional regression of normalized trade-time volatilities from one trading condition on those from another trading condition must equal one. We then show that this prediction is also implied by an intuitive "no speculative arbitrage" condition that the expected profit from investigating one asset rather than another must be the same.

We obtain weak support for bet invariance when using realized spot volatilities as inputs. Most slope coefficients are economically and statistically less than one. We also find evidence of significant noise in the estimates when the independent variable is associated with high trading activity. When we instead use expected normalized volatilities, based on the average of the previous 24 months, to proxy for expected volatility, we find support for the theory. This establishes that departures are driven by unexpected cross-sectional variations in volatility that bias coefficient estimates downward. However, we do not find supporting evidence from one time of day to another. In sum, we find strong evidence of the role of speculative trading in equating *expected* profit opportunities to speculators across assets and trading conditions, but not across time. Our findings shed light on the systematic cross-stock patterns of volatility at intraday frequencies.

6 References

Andersen, T.G., O. Bondarenko, A.S. Kyle, and A.A. Obizhaeva (2016). Intraday trading invariance in the E-mini S&P 500 futures market. Working paper.

Andersen, T.G., T. Bollerslev, and F.X. Diebold (2010). Parametric and nonparametric volatility measurement. In *Handbook of Financial Econometrics, Volume 1—Tools and Techniques.*, ed. Yacine Ait-Sahalia and Lars Peter Hansen, 67–137. North Holland.

Ang, A., R.J. Hodrick, Y. Xing, and X. Zhang (2006). The cross-section of volatility and expected returns. Journal of Finance 61, 259–299.

Angel, J., L. Harris, C. Spatt (2011). Equity Trading in the 21st Century. Quarterly Journal of Finance 1, 1–53.

Axtell, R.L. (2001). Zipf distribution of US firm sizes. Science, 293(5536), 1818–1820.

Bakshi, G. and N. Kapadia (2003). Delta-hedged gains and the negative market volatility risk premium. Review of Financial Studies 16, 527–566.

Barardehi, Y.H. and D. Bernhardt (2017). Intraday dynamics of trade-time outcomes. Working paper.

Barardehi, Y.H., D. Bernhardt, and R.J. Davies (2016). Trade-time measures of liquidity. forthcoming. Review of Financial Studies.

Baruch, S. and L.R. Glosten (2013). Fleeting orders. Working paper.

Brennan, M.J., T. Chordia, A. Subrahmanyan (1998). Alternative factor specifications, security characteristics, and the cross-section of expected stock returns. Journal of Financial Economics 49, 345–373

Budish, E., P. Cramton, and J. Shim (2015). The high-frequency trading arms race: frequent batch auctions as a market design response. Quarterly Journal of Economics 130(4), 1547–1621.

Cao, J. and B. Han (2013). Cross section of option returns and idiosyncratic stock volatility. Journal of Financial Economics 108, 231-249.

Conrad, J., S. Wahal, J. Xing (2015). High frequency quoting, trading, and the efficiency of prices. Journal of Financial Economics 116(2), 271–291.

Dufour, A., R.F. Engle (2000). Time and the price impact of a trade. Journal of Finance 55(6), 2467–2598.

Easley D., M.M. Lopez de Pardo, M. O'Hara (2012). Flow toxicity and liquidity in a high frequency world. Review of Financial studies 25(5), 1457–1493.

Frisch, R. Waugh, F.V. (1933). Partial Time Regressions as Compared with Individual Trends. Econometrica 1(4), 387–401.

Gabaix, X., A. Landier. (2008). Why has CEO pay increased so much?. The Quarterly Journal of Economics, 123(1), 49–100.

Gai, J., C. Yao, and M. Ye (2013). The externalities of hig-frequency trading. Working paper.

Glosten, L.R. (1994). Is the electronic open limit order book inevitable? Journal of Finance 47(4), 1127–1161.

Gouriėroux, C., J. Jasiak, and G. Le Fol (1999). Intra-day market activity. Journal of Financial Markets (2), 193–226.

Hasbrouck, J. (1999). Trading fast and slow: security market events in real time. Working paper, New York University.

Hasbrouck, J. and G. Saar (2013). Low-latency trading. Journal of Financial Markets (16), 646–679.

Hendershott, T., R. Riordan (2009). Algorithmic trading and information. Working paper, University of California Berkley.

Holden, C.W., S.E. Jacobsen (2014). Liquidity measurement problems in fast, competitive markets: Expensive and cheap solutions. Journal of Finance 69(4), 1747–1785.

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

Kyle, A.S. and A. Obizhaeva (2016). Market microstructure Invariance: empirical hypotheses. Econometrica 84(4), 345–1404.

Ijiri, Y. and H. Simon (1977). Skew distributions and the sizes of business firms (Vol. 24). North Holland.

Lehmann, B.N. (1990). Residual risk revisited. Journal of Econometrics 45, 71–97.

O'Hara, M., C. Yao, and M. Ye (2014). What's not there: odd-lots and market data. Journal of

Finance 69(5): 2199–2236.

Tinic, S.M. and R.R. West (1986). Risk, return and equilibrium: A revisit. Journal of Political Economy 94, 126–147.

Upson, J., H. Johnson, and T.H. McInish (2015). Orders versus trades on the consolidated tape. Working paper.

Wood, R.A. , T.H. McInish, and J.K. Ord (1985). An investigation of transaction data for NYSE stocks. Journal of Finance 60, 723–739.