

**Centre for
Economic
and Financial
Research
at
New Economic
School**



March 2016

Intraday Trading Invariance in the E-mini S&P 500 Futures Market

Torben G. Andersen
Oleg Bondarenko
Albert S. Kyle
Anna A. Obizhaeva

Working Paper No 229

CEFIR / NES Working Paper series

Intraday Trading Invariance in the E-mini S&P 500 Futures Market

TORBEN G. ANDERSEN, OLEG BONDARENKO
ALBERT S. KYLE AND ANNA A. OBIZHAEVA*

First Draft: July 2, 2014

This Draft: March 15, 2016

The intraday trading patterns in the E-mini S&P 500 futures contract between January 2008 and November 2011 are consistent with the following invariance relationship: The return variation per transaction is log-linearly related to trade size, with a slope coefficient of -2 . This association applies both across the pronounced intraday diurnal pattern and across days in the time series. The documented factor of proportionality deviates sharply from prior hypotheses relating volatility to transactions count or trading volume. Intraday trading invariance is motivated a priori by the intuition that market microstructure invariance, introduced by Kyle and Obizhaeva (2016c) to explain bets at low frequencies, also applies to transactions over high intraday frequencies.

Keywords: market microstructure, invariance, bets, high-frequency trading, liquidity, volatility, volume, business time, time series, intraday patterns.

* Andersen: Kellogg School of Management, Northwestern University, 2001 Sheridan Road, Evanston, IL 60208, USA, NBER, and CREATES, t-andersen@kellogg.northwestern.edu. Bondarenko: Department of Finance (MC 168), University of Illinois at Chicago, 601 S. Morgan St., Chicago, IL 60607, USA, olegb@uic.edu. Kyle: Robert H. Smith School of Business, University of Maryland, College Park, MD 20742, USA, akyle@rhsmith.umd.edu. Obizhaeva: New Economic School, 100A Novaya Street, Skolkovo, Moscow, 143026, Russia, aobizhaeva@nes.ru. We are grateful to the CME Group for access to data from the CME DataMine system. Andersen acknowledges support from CREATES, Center for Research in Econometric Analysis of Time Series (DNR78), funded by the Danish National Research Foundation. Kyle works as a director for an asset management company which trades international stocks and has also worked as a

I. Introduction

An extensive literature has documented systematic intraday variation in market activity variables such as trading volume, number of transactions, and return volatility. This variation is highly interdependent, e.g., an unusually large trading volume is typically associated with high return volatility. At least qualitatively, such relationships are common across time, asset classes, and market structures. Thus, it is natural to view this type of covariation as a universal property that arises endogenously from the interaction of trading strategies, the daily rhythm of news releases, and business activity. Nonetheless, little is known about any specific quantitative association between this set of activity variables at the intraday level. In this paper, we examine whether the number of trades, the size of individual trades, and return volatility have a precise quantitative relationship which remains invariant both within and across trading days. In particular, we explore whether these intraday interactions are consistent with predictions that arise from the invariance principle developed in Kyle and Obizhaeva (2016*c*).

Unfortunately, accurate identification of the number and size of trades by active investors is difficult to infer from transaction data published by most exchanges. First, in electronic order book markets, it is typical to record large trades, which cross several standing limit orders, as separate trades. Thus, a marketable order for, say, 12 units may be executed against four different limit orders with sizes of 2, 1, 5, and 4 units, respectively. This will appear as four separate trades at the identical price, reflecting the execution of four (passive) limit orders, rather than as a single (active) order. In our motivating framework, however, the endogenous choice of transaction intensity and size by active investors is critical, rendering this way of recording trades incompatible with the underlying theory. Second, many markets are fragmented, with trading taking place simultaneously at several distinct venues. This raises further questions about proper sequencing and integration of the trading activities across market segments.

To avoid these complications, we focus our empirical study on a single market, the S&P 500 E-mini futures contract traded on the Globex electronic platform of the CME Group. This market has a number of advantages for our purposes. First, it is an important and extremely liquid market; it is regarded as the prime venue for price discovery in the U.S. equity markets, and the daily turnover is among the top two globally for exchange-traded assets. Second, it is a fully integrated market, as the E-mini contract trades only through the Globex system; this greatly simplifies identifying the active trades and sequencing them in correct order. Third, the

consultant for various U.S. government agencies, including the SEC and CFTC. We appreciate comments from participants at the Workshop on Recent Advances in High-Frequency Statistics, Humboldt University, Berlin; the conference Market Microstructure—Confronting Many Viewpoints, Paris; the High-Frequency Financial Data Conference, Montréal; the 7th R/Finance Conference, University of Illinois, Chicago; the European Summer Symposium in Financial Markets, Gerzensee, Switzerland; the 8th Annual Conference of the Society for Financial Econometrics, Aarhus, Denmark; and the 11th World Congress of the Econometric Society, Montréal.

market operates almost continuously during weekdays, producing a huge degree of variation in the activity patterns across the Asian, European and American trading regions; this provides us with excellent statistical power to identify the intraday interdependencies among the market activity variables. Fourth, the data from the CME Group directly identifies the transaction size from the perspective of the active participant; consequently, for the example contemplated above, we observe a single trade of 12 contracts. Fifth, the large tick size increases the number of contracts available at the best bid and offer; hence, the market depth is usually sufficiently large that almost all orders, including entire meta-orders for hundreds of contracts or more, may be consummated at once, rather than being broken into smaller pieces and executed over time.

Figure 1 depicts the average trading volume V , return volatility σ , number of transactions N , and trade size Q across the 24-hour cycle of trading in the E-mini futures market. The series are generated from tick data and show the average values within consecutive one-minute intervals for January 2008 to November 2011. The exact construction of the series is detailed in Section III. The figure displays calendar time on the horizontal axis with the three trading regimes separated by dashed vertical lines. The first regime spans -7:00 (17:00 of the prior day) to 2:00 CT (Chicago Time), covering the period before and during regular trading in Asia; the second regime encompasses 2:00 to 8:30 CT, roughly corresponding to standard European trading hours; and the third regime covers 8:30 to 15:15 CT, matching active American trading hours. The combined 22 hour and 15 minute segment represents continuous trading in the Globex market.

The extreme variation in the intraday, or diurnal, pattern across the regions is evident from the figure. Moreover, the pattern is clearly distinct for the separate activity variables. The trading volume increases more than fifty-fold and the transaction count grows twenty five-fold when trading moves from Asian to American hours; in contrast, volatility and trade size only roughly double.

Our primary empirical specification is motivated by the market microstructure invariance hypothesis of Kyle and Obizhaeva (2016*c*), stipulating a particular interdependence between the variation in average trade size, transaction frequency, and return volatility. Invariance is based on the intuition that markets for different financial assets operate similarly when viewed from the perspective of an appropriate security-specific “business-time” clock, which ticks at a rate proportional to the expected arrival of “bets” or meta-orders. Bets represent the sum of unidirectional trades spread across periods which might be minutes, hours, days, or even weeks. The theory implies that the arrival rate of bets is proportional to the $2/3$ power of the product of expected dollar volume and return volatility. Equivalently, the return variance per unit of business time is proportional to the inverse of the squared bet size. Using portfolio transition orders as proxies for bets, Kyle and Obizhaeva (2016*c*) show that variation in the average bet size across stocks with different levels of trading activity conforms to the predictions of invariance.

This paper examines empirically whether the invariance principle can be ex-




FIGURE 1. THE FIGURE SHOWS AVERAGES ACROSS ALL DAYS FOR CONTRACT VOLUME V_t (PER MINUTE), ANNUALIZED REALIZED VOLATILITY σ_t , TRADE SIZE Q_t , AND NUMBER OF TRADES N_t . THE AVERAGES ARE COMPUTED AT A GRANULARITY OF ONE MINUTE. THE DASHED VERTICAL LINES SEPARATE THE THREE TRADING REGIMES CORRESPONDING TO TRADING HOURS IN ASIA, EUROPE, AND AMERICA.

tended to rationalize the variation in the number and average size of transactions within the trading day.¹ Specifically, we develop alternative log-linear invariance representations and test whether they hold up over short intraday intervals like one minute. Since we measure variables at very high frequencies, we refer to our hypothesis as the “intraday trading invariance hypothesis” or, more simply, as “intraday invariance.” We confirm that the return variance per transaction is approximately inversely proportional to the power of average trade size through estimates of an exponent close to the theoretical value of -2 . We also confirm that the number of transactions is proportional to the power of the product of expected dollar volume and return volatility with an exponent close to $2/3$.

Intraday invariance differs significantly from earlier theories describing how volume and volatility vary over time. We demonstrate that invariance provides a substantially better characterization of the interactions among key market activity variables than alternative theories motivated by the earlier work of Mandelbrot

¹Work considering variation in trade size include Brennan and Subrahmanyam (1998) for the cross-section as well as Alexander and Peterson (2007) and Moulton (2005) for the time series.

and Taylor (1967) or Clark (1973), using, respectively, the number of transactions or volume as the directing process for volatility.

We also examine separately how intraday invariance performs during times of extreme market activity. We find that it holds up well following the release of scheduled macroeconomic announcements. During the flash crash of May 6, 2010, intraday invariance was operative while prices dived, but it broke down at the very end of the crash and was inoperative during the initial phase of the recovery. Intraday invariance also breaks down for a short period around the times when regional market centers open and close.

These deviations during market transitions suggest that intraday invariance may be tied to inherent features of properly functioning markets within a stable environment. The instability may be induced by changes in the order-shredding algorithms, intermediation strategies, and order-matching mechanisms as activity shifts from one global trading region to another. As such, these regular departures from intraday invariance do not speak to the validity of the motivating theory based on the market microstructure invariance hypothesis of Kyle and Obizhaeva (2016c), which concerns the arrival rate of bets. Instead, we conjecture that violations of intraday trading invariance occurring *inside* a well-defined regional trading segment signify a breakdown in the normal state of the market.

II. Invariance Hypotheses

Kyle and Obizhaeva (2016c) develop market microstructure invariance by applying an invariance principle to bets. Intraday invariance is based on applying an analogous principle to the intraday trading process.

Below, we review the market microstructure invariance hypothesis in Section II.A; provide an extension of the theory to an intraday setting in Section II.B; detail our methodology for generating empirical proxies for the unobserved market activity variables in Section II.C; explain how we alleviate the measurement errors in Section II.D; clarify how intraday trading invariance differs from alternative hypotheses relating volume, volatility, and trading activity in Section II.E; and develop a nested regression-based test of the alternative theories for the interaction among the market activity variables in Section II.F.

A. Review of Market Microstructure Invariance

The market microstructure invariance hypothesis of Kyle and Obizhaeva (2016c) is based on the intuition that markets for different financial assets operate at different business-time scales. Business time passes at a rate proportional to the frequency with which bets arrive into the market. A bet represents a decision by an investor to change his risk exposure to the asset. In modern electronic markets, large bets are typically shredded into numerous smaller orders which are executed strategically over time to minimize price impact costs. For example, if a trader decides to purchase 1,000 futures contracts and implements this decision by

making 100 unidirectional purchases during one hour, the first purchase represents the arrival of the bet while the subsequent 99 do not represent new bets.

Kyle and Obizhaeva (2016*b*) describe a theoretical model in which return volatility results from the linear price impact of bets, noise traders bet on uninformative signals, and market makers take the other side of trades with informed and noise traders. When noise traders place more bets, profit opportunities increase and informed traders place more bets. The rate at which informed traders place bets reflects a trade-off between the cost of acquiring private information (an exogenous constant) and the expected trading profits net of market impact costs (decreasing in the number of informed traders). In equilibrium, they are equated due to a free-entry condition. This further leads to the market microstructure invariance hypothesis, which states that the dollar risk transferred by a bet is invariant when measured in units of business time, directed by the rate of bet arrival.²

To express this invariance relationship formally, we need some notation. We use a tilde to signify random variables, while we refer to the corresponding expected values by omitting the tilde. We let P denote the asset price (dollars per contract), \tilde{Q}_B denote the random unsigned bet size (number of contracts), σ^2 denote the percentage return variation (per unit time), and N_B denote the arrival rate of bets (number per unit of time). The bet arrival rate N_B measures the speed with which business time passes. Formally, define the invariant \tilde{I} by,

$$(1) \quad \tilde{I} := P \cdot \tilde{Q}_B \cdot \frac{\sigma}{N_B^{1/2}}.$$

The invariance hypothesis predicts that active traders scale their (random) bet size \tilde{Q}_B in response to market conditions ($P \cdot \sigma \cdot N_B^{-1/2}$) so that \tilde{I} is independently and identically distributed across assets and time. Bets arrive at the expected rate N_B , and the size of each bet \tilde{Q}_B is then drawn randomly but consistent with the invariance of the \tilde{I} distribution.

Empirical analysis of invariance is complicated by the fact that business time—the bet arrival rate N_B —is not easily observable at any point in time. Identifying bets requires assigning individual transactions to specific traders and aggregating them into bets. Thus, inherently, one ends up relying on time-aggregation to estimate the intensity of bet arrivals. Likewise, while one may readily devise estimators for realized return volatility over intervals of non-trivial length, it is difficult to estimate expected volatility at any point in time—the so-called spot volatility. As for trading intensity, volatility estimation involves aggregation and averaging. An additional complication is that both trading intensity and return volatility are highly variable across the trading day.

Throughout the remainder of this section, we assume, for expositional simplicity, that bet volume equals total volume, e.g., one side of every trade is a bet and the

²There are additional theoretical invariance relations linking, e.g., trading activity to bid-ask spreads, price impact, and price efficiency, but we do not examine any of these in this paper.

other side is an intermediation or market-maker trade.³ Thus, letting $Q_B = E[\tilde{Q}_B]$ denote the expected bet size and V the expected volume (contracts per unit of time), we have $V = N_B \cdot Q_B$.

Now, defining expected “trading activity” W as the standard deviation of the dollar mark-to-market gains or losses on the expected volume per unit of calendar time, we obtain $W = P \cdot V \cdot \sigma$. Intuitively, the quantity W measures the “expected total risk transfer” per unit of calendar time.

Invariance implies that the rate at which business time passes is proportional to the 2/3 power of trading activity W . To see this, exploit the definitions of W and V to write,

$$(2) \quad \tilde{I} = \frac{P \cdot \tilde{Q}_B \cdot \sigma}{N_B^{1/2}} = \frac{\tilde{Q}_B}{Q_B} \cdot \frac{P \cdot V \cdot \sigma}{N_B^{3/2}} = \frac{\tilde{Q}_B}{Q_B} \cdot \frac{W}{N_B^{3/2}}.$$

Let $I = E[\tilde{I}]$ denote the mean of the invariant distribution. We obtain,

$$(3) \quad I = \frac{P \cdot Q_B \cdot \sigma}{N_B^{1/2}} = \frac{W}{N_B^{3/2}}.$$

Since \tilde{I} has an invariant distribution, its mean I is constant and independent of W or N_B . It follows that N_B is proportional to $W^{2/3}$.

Our empirical work focuses on alternative versions of the linear representation that arises from taking the logarithm on both sides in equation (1). Define $p := \log P$, $\tilde{q}_B := \log \tilde{Q}_B$, $s := \log \sigma^2$, $n_B := \log N_B$, and $q_B := E[\tilde{q}_B]$. Then,

$$(4) \quad E \left[\log \tilde{I} \right] = p + q_B + \frac{1}{2} s - \frac{1}{2} n_B.$$

Note that the empirical content of equation (4) differs slightly from equation (3) because, due to Jensen’s inequality, the means of the logs of \tilde{Q}_B and \tilde{I} are smaller than the logs of the means. Nonetheless, equations (3) and (4) are both valid representations, since invariance implies the *distribution* of \tilde{I} is invariant.

The hypothesis that the right-hand-side of equation (4) is constant may be expressed in various ways. Here, we provide a specific representation that has served as the basis for prior empirical work, but we later develop other regression specifications for testing invariance across the intraday trading patterns.

Defining the log of trading activity as

$$(5) \quad w = \log W = p + v + \frac{1}{2} s,$$

we readily obtain the proportionality relation,

$$(6) \quad n_B = c + \frac{2}{3} \cdot w,$$

³Similar predictions apply in general if total volume is proportional to bet volume.

where we let c represent a generic constant whose value typically differs across equations. Since $v = n_B + q_B$, the same relationship may equivalently be written with the log bet-size q_B on the left-hand-side,

$$(7) \quad q_B = c + \frac{1}{3} \cdot w - \left[p + \frac{1}{2} \cdot s \right].$$

Equations (6) and (7) formalize the intuition that, for fixed s and p , as the trading activity w grows, 2/3 of the increase stems from the arrival rate of bets (speed of business time) and 1/3 from the magnitude of bets. This relationship between the size and number of bets lies at the heart of market microstructure invariance.

Several empirical hypotheses inspired by invariance find support in the data. Kyle and Obizhaeva (2016c) document invariance relationships for the size distributions of portfolio transition orders. Kyle, Obizhaeva and Tuzun (2016) document an invariance relationship for the size distribution of transactions in the Trade and Quote data set (TAQ). Kyle et al. (2014) study invariance relations for the number of monthly news articles. Bae et al. (2014) discuss an invariance relationship for the number of buy-sell switching points in the South Korean equity market.

We now turn to an exploration of intraday trading invariance for the high-frequency patterns in the E-mini S&P 500 futures market. Our tests are not a direct implication of an underlying invariance hypothesis. Instead, like Kyle, Obizhaeva and Tuzun (2016), they are based on extrapolating the invariance hypothesis from bets to intraday transactions.

B. The Intraday Trading Invariance Hypothesis

“Intraday trading invariance” is the hypothesis that the invariance relationships apply not only to bets but also to trading variables observed at intraday intervals. One way to ensure that invariance relations for bets may be extrapolated to intraday data is to assume that the average number of transactions per bet is also invariant across assets and time. This is merely a sufficient condition, however, and we hypothesize that intraday invariance applies much more generally.

To define the intraday-invariance hypothesis formally, we introduce notation referring to specific intraday calendar intervals. The sample begins at time 0 and contains D trading days, each comprised of T intraday intervals of length $\Delta t = 1/T$. Thus, the full sample covers $[0, D]$ and contains the non-overlapping intervals $\tau = 1, \dots, D \cdot T$.

For any interval τ , the average transaction price is denoted P_τ , the average number of units traded per transaction is Q_τ , the conditionally expected return volatility given information prior to time τ is σ_τ , and the conditionally expected transaction rate (number of trades per unit time) is N_τ .

Intraday Trading Invariance: *Intraday trading invariance, or simply intraday invariance, implies that \tilde{I}_τ , defined as,*

$$(8) \quad \tilde{I}_\tau := \frac{P_\tau \cdot \tilde{Q}_\tau \cdot \sigma_\tau}{N_\tau^{1/2}}.$$

has an invariant distribution for all intervals τ .

The notation reflects the fact that P_τ is observed, while both σ_τ and N_τ represent (unobserved) conditional expectations at time τ , characterizing the current trading environment. These variables are viewed as known to the agents implementing the transactions stemming from bet arrivals. The hypothesis then implies that traders endogenously modify their transaction size based on perceived market conditions, so that the distribution for \tilde{Q}_τ satisfies the invariance principle (8).

It follows from equation (8) that $\tilde{i}_\tau = \log \tilde{I}_\tau$ also constitutes an i.i.d. sequence. Taking expectations conditional on the information available prior to interval τ implies that a specific linear combination of the expected logarithmic trading variables equals an invariant constant. In the notation established for equation (4), we have, for all τ with strictly positive trading volume,

$$(9) \quad E[\tilde{i}_\tau] = p_\tau + q_\tau + \frac{1}{2}s_\tau - \frac{1}{2}n_\tau.$$

By analogy with equation (6), we obtain,

$$(10) \quad n_\tau = c + \frac{2}{3} \cdot w_\tau.$$

Intraday invariance is much easier to test empirically than the corresponding invariance hypothesis for bets. First, while bets are difficult to observe, the size and number of trades are public information which can be observed in time-stamped tick data. Second, in today's twenty-four-hour global markets, huge fluctuations in trading volume and volatility are observed across the Asian, European, and American trading regimes. These features provide us with good statistical power to detect whether the intraday fluctuations for trade size Q_τ and expected number of trades N_τ are related to the intraday return volatility σ_τ in a manner consistent with the proportions implied by the invariance relation.

Nevertheless, since the conditionally expected trade size, transaction rate and volatility are not directly observed and vary substantially, both over time and across the diurnal cycle, it is not possible to exploit equation (9) directly in a statistical test. First, we must design suitable estimators for the latent variables. For that purpose, we assume that the variation in the conditional expectation is negligible over each short intraday interval τ and then apply standard techniques to obtain unbiased estimators for the level of each trading variable, interval by interval. Such estimators are inevitably quite noisy and the further transformation into quantities suitable for inclusion in tests of our basic linear representation involves a few subtleties, which we discuss below.

C. Methodology for High-Frequency Measurement

The intraday-invariance relation involves a set of non-negative random variables that are persistent, yet also contain large idiosyncratic innovations. Such variables are not compatible with standard time-series models featuring Gaussian errors.

Instead, we turn to the so-called multiplicative error model (MEM) approach. This refers to a broad family of models, including all standard GARCH and stochastic volatility specifications. The framework has a long history within the modeling of intraday market data. Originally, it was used to capture high-frequency return dynamics through the interaction of persistent daily volatility factors and pronounced intraday diurnal features, see Andersen and Bollerslev (1997, 1998). It was cast as a more general modeling framework and labeled MEM by Engle (2002), with applications to both volatility and trade durations. Finally, it has been used successfully in modeling the dynamics of high-frequency trading volume in Brownlees, Cipollini and Gallo (2011, 2012).

We let \tilde{Y}_τ denote the value of a strictly positive random variable (representing V , N , or σ^2), and then stipulate that the dynamics of \tilde{Y}_τ is given by,

$$(11) \quad \tilde{Y}_\tau = Y_\tau \cdot \tilde{U}_\tau,$$

where \tilde{U}_τ is strictly positive, i.i.d., with the mean of one and variance of $\sigma_{\tilde{U}}^2$. Equation (11) is a MEM model with time $\tau - 1$ conditional expectation $E_{\tau-1}[\tilde{Y}_\tau] = Y_\tau$ and $Var_{\tau-1}[\tilde{Y}_\tau] = Y_\tau^2 \cdot \sigma_{\tilde{U}}^2$. The conditional dependence structure is captured by the positive mean, Y_τ , which remains unspecified and can be highly complex, incorporating interdaily dependencies, diurnal patterns, and short-run serial correlation.

The critical MEM feature is the identically distributed multiplicative error structure. Even if we impose no direct distributional assumptions on these errors, the representation imposes a strong restriction on the system, as is evident from the specific form of the conditional variance. However, in the case of a continuous return volatility process, the spot variance of σ^2 takes this exact form.⁴ Likewise, for trading activity variables like N and V , it has been confirmed that the scale of the innovations are approximately proportional to the (conditional) mean, see, e.g., Andersen (1996) and Brownlees, Cipollini and Gallo (2011).

In this setting, we may construct simple unbiased estimators for σ_τ^2 , V_τ and N_τ , subject to standard regularity conditions, as long as the interval $[\tau - \Delta t, \tau]$ is sufficiently short so that the intra-interval variation in the expected value is negligible. For volatility, we cumulate very high-frequency squared returns from within the interval τ while, for the transaction variables, we may rely on the actual observed volume and transaction counts over the interval. Of course, sensible measures require a minimum of market activity, so we discard the interval if any of these three estimators is zero. Hence, our subsequent tests for intraday trading invariance only involve periods with a non-trivial degree of market activity.

⁴This fact underlies the well-known result that the asymptotic variance of the realized volatility estimator is proportional to the integrated quarticity.

Taking the logarithm of equation (11) and letting lower case letters denote the corresponding logarithmic values, so, e.g., $E_{\tau-1}[\tilde{y}_\tau] = y_\tau$, we have,

$$(12) \quad \tilde{y}_\tau = y_\tau + \tilde{u}_\tau,$$

where \tilde{u}_τ are i.i.d. random variables with mean $\mathbf{E}[\log \tilde{U}_\tau]$ and finite variance. The “innovation” term \tilde{u}_τ is not centered, but has a negative mean, due to Jensen’s inequality, $E[\tilde{u}_\tau] < \log \left(E[\tilde{U}_\tau] \right) = 0$. Importantly, however, the mean of \tilde{u}_τ is constant across τ , $E[\tilde{u}_\tau] = c < 0$. Hence, equation (12) implies,

$$(13) \quad \tilde{y}_\tau = y_\tau + c + \tilde{\epsilon}_\tau,$$

where $\tilde{\epsilon}_\tau$ are i.i.d. random variables with zero mean and finite variance, while c denotes a generic constant which will differ across the activity variables.

Equation (13) provides guidance for the construction of estimators amenable for inclusion in log-linear regression-based tests. We construct unbiased estimators for the variables themselves, which typically may be obtained through suitable sample averages over a short interval. The logarithm of this sample average yields an estimator with an invariant bias which may be absorbed into the constant term in regression tests. Such estimators are, however, very noisy due to the presence of the sizable error term, $\tilde{\epsilon}_\tau$, in equation (13). Below, we pursue the strategy of mitigating the impact of measurement errors by averaging the estimator across a large number of separate intervals, thus diversifying these errors.

D. Alternative Aggregation Schemes

Our high-frequency measures of the expected activity variables, σ_τ , N_τ , and V_τ are invariably noisy, so testing for intraday trading invariance requires effective error mitigation. This section develops alternative schemes to mitigate the impact of measurement errors and allow for natural interpretation of the test results.

To facilitate the exposition, it is useful with notation that directly identifies the trading day and intraday time period associated with a given interval, τ . Specifically, we denote the set of all trading days $\mathcal{D} = \{1, \dots, D\}$ and the set of intraday intervals $\mathcal{T} = \{1, \dots, T\}$. Our full sample, covering $[0, D]$, is thus comprised of the set of non-overlapping intervals $\mathcal{D} \times \mathcal{T}$. We now conveniently refer to any interval via the double-index (d, t) , referring to interval t on day d , which corresponds to interval $\tau = (d - 1) \cdot T + t$ in our single-index notation.

Our strategy is to aggregate the high-frequency observations across distinct subsets of intervals to mitigate the measurement errors, while retaining a wide dispersion in the (average) value of the activity variables across these categories. The latter feature ensures that we have good test power when confronting the intraday invariance predictions with the observed variation across subsets.

We denote a generic subset of trading days and intraday intervals by $\mathcal{D}_k \subseteq \mathcal{D}$ and $\mathcal{T}_j \subseteq \mathcal{T}$, respectively. We average the random variable \tilde{y}_{dt} across all intervals

in $\mathcal{D}_k \times \mathcal{T}_j$. Letting $|\mathcal{A}|$ denote the number of elements in the set \mathcal{A} , we construct,

$$(14) \quad \tilde{y}_{k,j} = \frac{1}{|\mathcal{D}_k|} \frac{1}{|\mathcal{T}_j|} \cdot \sum_{d \in \mathcal{D}_k, t \in \mathcal{T}_j} \tilde{y}_{dt},$$

where index k corresponds to aggregation in the time-series dimension and index j to aggregation across the intraday pattern. Equations (13) and (14) imply,

$$(15) \quad \tilde{y}_{k,j} \approx y_{k,j} + c,$$

with

$$(16) \quad y_{k,j} = \frac{1}{|\mathcal{D}_k|} \frac{1}{|\mathcal{T}_j|} \cdot \sum_{d \in \mathcal{D}_k, t \in \mathcal{T}_j} y_{dt} \quad \text{and} \quad \frac{1}{|\mathcal{D}_k|} \frac{1}{|\mathcal{T}_j|} \cdot \sum_{d \in \mathcal{D}_k, t \in \mathcal{T}_j} \tilde{\epsilon}_{dt} \approx 0$$

following from the law of large numbers, provided $|\mathcal{D}_k| \cdot |\mathcal{T}_j|$ is large.

We simplify the notation associated with our most common schemes. One approach fixes a given time of day and averages the observations for this interval across all days, i.e., we construct the subsets $\mathcal{D} \times \{t\}$, and define,

$$(17) \quad \tilde{y}_t = \frac{1}{|\mathcal{D}|} \cdot \sum_{d \in \mathcal{D}} \tilde{y}_{dt} = \frac{1}{D} \cdot \sum_{d=1}^D \tilde{y}_{dt}, \quad \text{for } t = 1, \dots, T,$$

which produces T separate observations across the intraday pattern.⁵

Analogously, we also pool observations across the Asian, European and American trading regimes within each trading day. In that case, the individual subset of intervals take the form $\{d\} \times \mathcal{T}_r$, with $r = 1, 2, 3$ representing the intraday intervals within each regime. The aggregate observations are then obtained as,

$$(18) \quad \tilde{y}_{d,r} = \frac{1}{|\mathcal{T}_r|} \cdot \sum_{t \in \mathcal{T}_r} \tilde{y}_{dt}, \quad \text{for } d = 1, \dots, D, \quad r = 1, 2, 3,$$

generating three distinct day-regime series, for a total of $D \cdot 3$ data points.

A similar example arises when we partition the regimes into smaller intervals, labelled bins, so the trading day comprises B distinct non-overlapping bins, \mathcal{T}_b . We denote a generic activity variable, averaged across the bins, $\tilde{y}_{d,b}$. Formally,

$$(19) \quad \tilde{y}_{d,b} = \frac{1}{|\mathcal{T}_b|} \cdot \sum_{t \in \mathcal{T}_b} \tilde{y}_{dt}, \quad \text{for } d = 1, \dots, D, \quad b = 1, \dots, B,$$

leading to a total of $D \cdot B$ data points.

⁵This involves some abuse of notation. Formally, we may refer to \tilde{y}_t as $\tilde{y}_{.,t}$, indicating averaging over the first index (and $\tilde{y}_{d,.}$ for averaging over the second index). Also, our prior variable \tilde{y}_r is distinct from \tilde{y}_t . The meaning should always be evident from the context.

E. Alternative Theoretical Hypotheses

There is a long history of theories describing and testing the relationship between trading activity and return volatility, dating back at least to Clark (1973). His work builds on the notion of subordination, or a stochastic business-time clock linked to trading activity, introduced into the modeling of financial returns by Mandelbrot and Taylor (1967).

Clark (1973) argues that increments to trading volume represent increments in business time, thus generating a close link between the expected daily volume and asset volatility. Subsequently, the “Mixture-of-Distributions Hypothesis” (MDH), developed in Tauchen and Pitts (1983) and Harris (1987), stipulates a linear relation, i.e., $\sigma_{dt}^2 \sim V_{dt}$, or equivalently in logs, $s_{dt} = c + v_{dt} = c + n_{dt} + q_{dt}$. Tauchen and Pitts (1983) note that the activity variable equally well may be the number of transactions as, in their model, the transaction size is assumed constant, rendering the two versions equivalent.

Subsequently, Jones, Kaul and Lipson (1994) find the number of trades to be better aligned with daily volatility, while Ané and Geman (2000) assert that intraday returns are i.i.d. Gaussian when normalized by the (stochastic) transaction count.⁶ The latter hypotheses suggest proportionality between the expected return variation and number of trades, i.e., $\sigma_{dt}^2 \sim N_{dt}$, or equivalently, $s_{dt} = c + n_{dt}$.

Both of these specifications, asserting a log-linear relation between volatility and trading activity, are derived via specializations of the basic MDH theory. We label them MDH-V and MDH-N, respectively.

We further note that, *if* the expected trade size is invariant over time, then the invariance relation (4) implies that the expected return variation is proportional to trading activity, $\sigma_{dt}^2 \sim N_{dt}$ and $\sigma_{dt}^2 \sim V_{dt}$. Thus, in the setting with constant trade size, intraday trading invariance is equivalent to both the MDH-V and MDH-N specifications. In contrast, if trade size varies with return volatility and trading activity, as suggested by Figure 1, equation (4) implies that the proportional relationship is systematically violated, and intraday trading invariance is incompatible with either of the two MDH specifications. We now turn toward empirical specifications that allow us to test the alternative hypotheses.

Using different aggregation schemes, we construct pooled series for the activity variables, which then serve as input to regression tests based on alternative log-linear representations. For example, following equation (17), we compute averages across all days, for given time-of-day t , to generate the series $\tilde{v}_t, \tilde{n}_t, \tilde{s}_t, \tilde{q}_t = \tilde{v}_t - \tilde{n}_t, \tilde{p}_t$ and \tilde{w}_t .⁷

From the perspective of practical regression analysis, however, we encounter a problem. In sharp contrast to the other quantities, the asset price has no systematic

⁶Andersen (1996), Bollerslev and Jubinski (1999), and Liesenfeld (2001) refine the relationship between trading variables and return volatility in specifications that extend the MDH model.

⁷The tilde above the variables indicates that our regression framework involves empirical proxies for the underlying conditional expectations. Hence, we explicitly acknowledge the presence of measurement or estimation error in these quantities.

intraday variation. For any intraday interval, the value of p_t , obtained as an average across all trading days, is effectively identical to the corresponding averaged price over any other interval.⁸ As a consequence, the expected variation in the price level is negligible, and the observed price variation largely reflects measurement error. This implies that inference regarding the association of the price level with the remaining variables will be extremely imprecise. Thus, for our regression tests, we ignore the price variable, folding p into the constant term c . This also results in specifications that nest intraday invariance and the two MDH alternatives.

Specifically, we alter the definition of w_τ in equation (5) to exclude p_τ ,

$$(20) \quad w_\tau = n_\tau + q_\tau + s_\tau/2 + c.$$

This modified measure captures the aggregate risk transferred in contracts, not in dollars. It enables us to express the MDH hypotheses through a 2/3-power representation comparable to the one developed for intraday invariance.

Given the sample averages of the activity variables, for given t , we readily obtain a regression representation of the intraday trading invariance relation (10),

$$(21) \quad \tilde{n}_t = c + \frac{2}{3} \cdot \tilde{w}_t + \tilde{u}_t.$$

Using the redefined w_t (with p_t treated as a constant), the MDH-V takes the form,

$$(22) \quad \tilde{n}_t = c + \frac{2}{3} \left[\tilde{w}_t - \frac{3}{2} \tilde{q}_t \right] + u_t^n.$$

Likewise, the MDH-N formulation may be restated as follows,

$$(23) \quad \tilde{n}_t = c + \frac{2}{3} [\tilde{w}_t - \tilde{q}_t] + u_t^n.$$

Obviously, the three parallel regression representations (21), (22), and (23) lend themselves to direct empirical comparison.

The regressions possess a rather unconventional feature. The composite regressor \tilde{w}_t , tautologically, covaries with the independent variable, as \tilde{n}_t is a component of \tilde{w}_t . This implies that the regressor is (positively) correlated with the error term, inducing an upward bias in the regression coefficient. It also follows that the R^2 of the regression is inflated, since a part of the idiosyncratic variation in the independent variable is captured, mechanically, through its presence in the regressor. Importantly, however, our error diversification, obtained by averaging across all days in the sample, should reduce the relative size of the error term drastically, thus greatly alleviating this endogeneity bias. Hence, it remains of interest to check, informally, whether the slope of the regression is roughly of the size predicted by the theories. In the following sections, we develop an alternative valid regression-based test for intraday trading invariance.

⁸In essence, this is a manifestation of the no-arbitrage principle, as significant and predictable intraday price level variation would generate simple profitable trading opportunities.

F. An Encompassing Regression Test

The three alternative hypotheses explored above each stipulate a specific relation between the market activity variables, \tilde{s}_t , \tilde{n}_t , and \tilde{q}_t . These variables are, empirically, observed with error, potentially creating problems for inference through the classical errors-in-variables effect. Our schemes designed to control and diversify the measurement errors mitigate these issues. We now seek an encompassing regression representation to allow for a direct comparison of the alternative hypothesis.

A natural approach is to include all variables in a multivariate regression specification. Unfortunately, this is problematic due to severe multicollinearity in the (relative) intraday variation of the series, as is also evident from Figure 1. The pairwise correlations of the activity variables are exceedingly high at 0.985 for n and s , 0.975 for n and q , and 0.936 for s and q . No matter what quantity is chosen as regressand, the extreme correlation among the two regressors will renders reliable estimation of the regression coefficients infeasible. Thus, to obtain a specification with power to discriminate among the hypotheses, we impose an a priori theoretical restriction which – importantly – is consistent with each of the three individual hypothesis. All theories imply that, in the multivariate log-linear regression involving all three variables, the log variance, s , and the log transaction count, n , enter with the identical coefficient.

Since the MDH-V and MDH-N hypotheses view volatility as evolving uniformly in business time, governed by their respective notion of trading activity, they naturally regard the variation in volatility as the quantity to be rationalized. Therefore, we focus on nested specifications for the alternative theories that place the return variance per transaction ($s - n$) on the left-side of the regression. It is readily verified that all three hypotheses may be expressed as follows,

$$(24) \quad \tilde{s}_t - \tilde{n}_t = c + \beta \cdot \tilde{q}_t + u_t^s.$$

MDH-V implies $\beta = 1$, MDH-N stipulates $\beta = 0$, and our intraday trading invariance hypothesis asserts $\beta = -2$.⁹ In other words, the theories offer starkly divergent predictions regarding the association between fluctuations in average trade size and return variation per transaction.

It is straightforward to conduct inference on the validity of different hypotheses within this regression setting, and we do so in Section IV.

III. The E-mini S&P 500 Futures Data

We exploit the best bid and offer (BBO) files for the E-mini S&P 500 futures contract from CME DataMine. These “top-of-the-book” files provide tick-by-tick information regarding the best quotes, order book depth, trade prices, and trade

⁹For intraday trading invariance, note that equation (21) implies $\tilde{n}_t = c + \frac{2}{3}(\tilde{n}_t + \tilde{q}_t + \frac{1}{2}\tilde{s}_t)$, and equation (24) with the value of $\beta = -2$ follows immediately.

sizes, time-stamped to the second. Since the contract is traded exclusively on the CME Globex electronic platform, the data contain all transactions executed during our sample, covering nearly four years from January 4, 2008, to November 4, 2011. The E-mini contract trades almost twenty-four hours a day, five days a week. The exchange trading hours are governed by Chicago Time (CT), so, henceforth, all references to time will be expressed in CT.

The procedure behind the recording of trades and trade sizes is critical for our empirical tests. When an executable order arrives, it is often matched with more than one limit order resting at the top of the book at the time of execution. During our sample period, the exchange reported all contracts traded at an identical price against an incoming order as a single combined transaction quantity. Thus, the trade size is reported from the perspective of the active party placing the executable order. This convention is consistent with the notion of trade size associated with the intraday invariance principle. Since most markets do not record transactions in this manner, it will typically be impossible to test for invariance relationships at the high-frequency level, as we do here, unless one is able to identify the active orders initiating the recorded trades.¹⁰

The notional value of the E-mini S&P 500 futures contract is \$50 times the value of the S&P 500 stock index, denominated in index points. The contract has a tick size of 0.25 index points (\$12.50), equivalent to approximately 2 basis points of notional value during our sample. It has four expiration months per year. We use data for the front month contract until it reaches eight days to expiration, at which point we switch into the next contract. This ensures that we use the most actively traded contract throughout our analysis.

Our analysis focuses on the uninterrupted trading session from 17:00 to 15:15 the following day. Thus, our trading week opens on Sunday at 17:00 and ends on Friday at 15:15. Daily trading in the Globex market is initiated by a batch auction at 17:00. The contracts traded during the opening auction are recorded during the first second of trading after 17:00. To generate a sample with a homogeneous trading protocol, we remove all transactions within the first second after 17:00 from our analysis. There is also trading on Monday through Thursday between 15:30 and 16:30; this hour, which involves another batch auction, is not open on the first trading day of the week (Sunday afternoon) and post-holidays. In order to avoid issues of non-homogeneity, we also exclude this fairly inactive trading hour from our analysis.

Since intraday trading invariance should hold for any exogenously selected subset of intervals, the exclusion of irregular trading segments does not affect the validity of our testing procedures. Finally, because of unusually low activity, we discard

¹⁰For example, if the exchange counts each limit order involved in trading as a distinct trade, the transaction count will reflect the limit order flow from the supply side that is intermediating trades. This will inflate the number of (active) transactions and shrink the trade size relative to the procedure used by the CME Group for the E-mini contract over our sample period.

one day from the sample, resulting in a total of $D = 969$ trading days.¹¹

As described in Section I, there are three trading regimes approximately corresponding to the trading hours in Asia, Europe, and America, namely 17:00–2:00, 2:00–8:30, and 8:30–15:15. In line with the notation established in Section II.D, we represent these regimes as non-overlapping sets \mathcal{T}_r , $r = 1, 2, 3$. They consist of 540, 390, and 405 one-minute intervals, respectively. Similarly, our sample covers 250, 252, 252, and 215 trading days for each of the four years 2008–2011, respectively. We let \mathcal{D}_y represent the trading days in year $y = 1, 2, 3, 4$, so the observations for regime r in year y are contained in the set $\mathcal{D}_y \times \mathcal{T}_r$.

Figure 1 in Section I depicts the average volume V_t , realized volatility σ_t , number of trades N_t , and average trade size Q_t for each minute of the trading day, with observations averaged across days in the sample in accordance with equation (17). The characterization of interdependencies among the pronounced intraday fluctuations in these trading activity variables is, of course, a central issue in this paper. We provide further details regarding the computation of each variable below.

Figure 2 plots the time-series variation in the activity variables, averaged across one-minute observations within each regime each day in accordance with equation (18). Consistent with the evidence from the extant literature, the series display a large degree of commonality. Trading volume, volatility, and the number of trades rise during the Société Générale scandal in January 2008 and the collapse of Bear Stearns in March 2008, soar during the financial crisis from September 2008 to February 2009, and spike during the flash crash in May 2010 as well as the second half of 2011. In contrast, the average trade size drops during the same periods. Finally, all variables appear slightly subdued toward the end of the calendar year.

There are also systematic differences across these time series. The relative scale of the variables differ sharply across regimes. The volume and number of transactions for American trading hours dwarf those observed during Asian trading hours, while volatility and trade size are more comparable across regimes. This is in line with Figure 1, where we observe similar stark divergences in the degree of variation across regions. Table 1 provides detailed summary statistics.

As an empirical matter, it is problematic to explore predictions of intraday trading invariance when trading costs and transaction sizes are distorted by market frictions. One striking feature of Table 1 is the fact that the average bid-ask spread for all three regimes barely exceeds the minimum value of 0.25 index points. The spread is binding and equals a single tick almost always, implying the tick size is “large.” Similarly, Figure 3 documents that a disproportionate number of trades across all three regimes in the the E-mini futures market involve a single contract, implying that the contract-size constraint also is binding.¹²

¹¹There are also four days (two in each of 2008 and 2009) where the entire Asian segment is missing due to a belated start to trading, and there are instances of abbreviated trading sessions for the European and American segments. A full discussion of the data used in our analysis is provided in Appendix VI.A.

¹²For early work on the impact of market frictions on the trading process, see, e.g., Harris (1994), Angel (1997), Goldstein and Kavajecz (2000), and Schultz (2000).

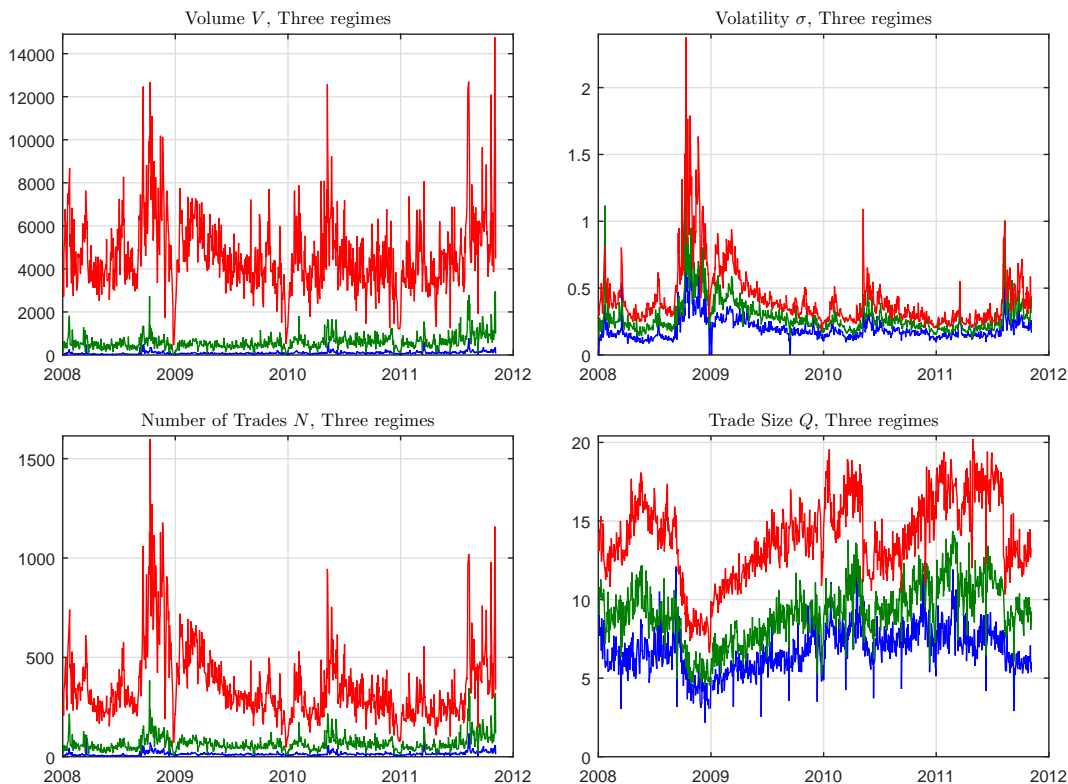


FIGURE 2. THIS FIGURE PLOTS THE TIME SERIES OF THE AVERAGE OF THE FOLLOWING ONE-MINUTE STATISTICS FOR EACH REGIME ON EACH TRADING DAY: VOLUME $V_{d,r}$ (PER MINUTE), VOLATILITY $\sigma_{d,r}$ (PER MINUTE, ANNUALIZED), NUMBER OF TRANSACTIONS $N_{d,r}$ (PER MINUTE), AND AVERAGE TRADE SIZE $Q_{d,r}$ (OVER ONE MINUTE), $r = 1, 2, 3$. THE ASIAN REGIME IS PLOTTED IN BLUE, THE EUROPEAN REGIME IS IN GREEN, AND THE AMERICAN REGIME IS IN RED. THE SAMPLE PERIOD IS JANUARY 4, 2008, TO NOVEMBER 4, 2011.

In contrast, the market is less restricted in terms of consummating large transactions. This is evident from the extreme size variation across trades depicted in Figure 4. While over half of the trades involve one or two contracts, there are, on average, a couple of transactions per minute during American trading hours exceeding 500 contracts. Table 1 shows that there is usually enough depth available at the best bid and ask to absorb trade sizes of hundreds of contracts. Large trades are the primary determinants of the average trade size, suggesting that this statistic is robust to distortions stemming from the minimum contract size. In this respect, the market provides a near ideal setting for exploring the interaction between trading intensity and volatility: it has an integrated electronic trading platform, a uniformly high degree of transaction activity, and a deep order book. The latter features are critical in enabling us to observe pronounced variation in the average trade size over time.¹³

We reiterate that the summary statistics in this section involve the observed

¹³Investigating the relationship between tick size and average trade size is an interesting issue which takes us beyond the scope of this paper.

TABLE 1—DESCRIPTIVE STATISTICS FOR THE E-MINI S&P 500 FUTURES

	Asia	Europe	America	Combined			
				Mean	Min	Max	Ratio
Volatility	0.16	0.25	0.40	0.26	0.12	0.81	6.5
Volume	95	601	4726	1647	51	32398	638
# Trades	14	67	360	134	9	1256	142
Notional Value, \$Mln	5	34	266	93	3	1814	626
Trade Size	5.9	8.4	13.3	8.9	4.6	28.5	6.2
Market Depth	54	265	984	398	35	3519	101
Bid-Ask Spread	26.5	25.7	25.1	25.9	25.1	28.0	1.1
Business Time	15.7	5.2	1.0	8.2	0.2	24.3	101

The statistics are reported separately for the three regimes. The last four columns show the mean, minimum, maximum, and maximum-to-minimum ratio for the entire day. Volatility is annualized. The volume, notional value, and number of trades are one-minute averages. Market depth is the average sum of the number of contracts at the best bid and ask. The bid-ask spread is measured in index points times 100. Business time is proportional to $W^{-2/3}$ and, for ease of comparison, it is normalized to be unity for America.




FIGURE 3. THE FIGURE PLOTS THE PROPORTION OF TRADES WITH SIZES $Q = 1$, $Q \leq 2$, AND $Q \leq 5$ CONTRACTS (LEFT PANEL) AND $Q \geq 200$, $Q \geq 100$, AND $Q \geq 50$ CONTRACTS (RIGHT PANEL) IN BLUE, GREEN, AND RED, RESPECTIVELY. THE PROPORTIONS ARE COMPUTED AT THE GRANULARITY OF ONE-MINUTE. THE DASHED VERTICAL LINES INDICATE THE THREE TRADING REGIMES.

one-minute trade and volatility data, while the theoretical relations underlying equations (21) and (22)-(24) concern the market's expected values for the logarithm of those variables. The one-minute sampling frequency represents a specific trade-off between bias and variance. On the one hand, if we sample at a lower frequency, we exploit more data in the computation of any given expectation and obtain a less noisy measure, reducing the variance of the estimator. On the other hand, as the interval grows, the underlying expectations will vary more substantially within the interval, inducing an increasing bias in the estimator due to the nonlinear relations we explore. By choosing a short one-minute interval, we deliberately seek to control this bias, while accepting noisy estimators. As discussed in Section II.D, it is therefore critical to mitigate the lack of precision by averaging

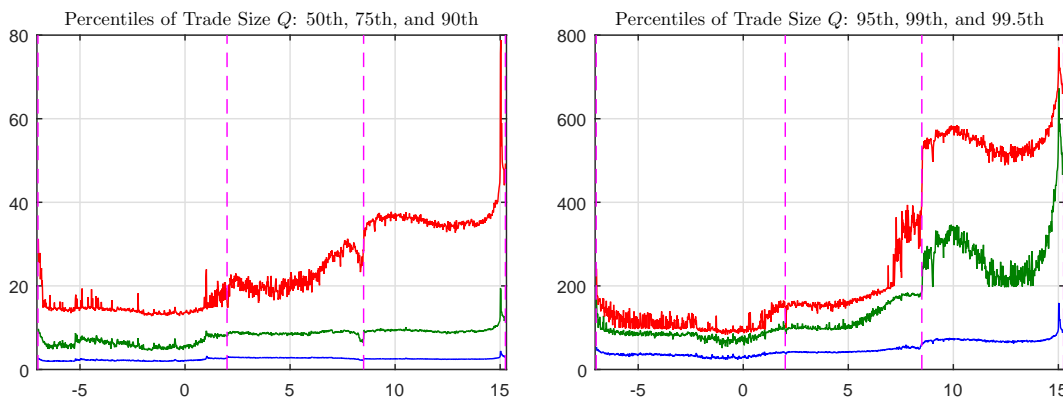


FIGURE 4. THE FIGURE PLOTS PERCENTILES OF THE TRADE SIZE Q . THE STATISTICS ARE COMPUTED AT THE GRANULARITY OF ONE-MINUTE. THE DASHED VERTICAL LINES SEPARATE THE THREE TRADING REGIMES.

the noisy log-transformed variables over a large set of intervals. For the trade size and number of transactions, this is straightforward, as the realized values are directly observed from the tick data. For the latent volatility, due to the liquidity of the E-mini contract, we use six consecutive squared ten-second returns to compute a realized volatility measure without any major detrimental effects from microstructure noise. Thus, we obtain a noisy, yet effectively unbiased, proxy for the true spot variance, and we then proceed as for the trade variables, i.e., we subsequently average the estimates across numerous one-minute intervals. This exploits the same error diversification principle that accounts for the accuracy of standard high-frequency realized-volatility estimators.¹⁴

Finally, we note that some observations are lost as we only test for invariance when the market displays meaningful activity. As mentioned previously, we exclude all intervals that feature zero volume or realized volatility. The fraction of omitted observations in the three regimes are 29.48%, 8.65% and 1.42%. One may worry that the elimination of data *conditional* on zero intra-interval price changes may induce a bias to the estimated volatilities used in the analysis. As a robustness check, we also conduct our tests at the substantially lower five-minute sampling frequency, which enables us to retain almost all observations. We now filter only 1.63%, 0.32% and 0.29 % of the minutes in the three regimes. Our results are qualitatively identical and quantitatively very similar, so we only present results for the one-minute sampling. The latter provides a more challenging setup and allows us to enhance the granularity of the tests around key transition points, when the characteristics of the trading process shift abruptly.

¹⁴See, e.g., Andersen, Bollerslev and Diebold (2010) or Ait-Sahalia and Jacod (2014) for a review of realized volatility estimation from high-frequency data. The daily realized volatility for equities is typically obtained from 78 five-minute or 390 one-minute squared returns over the trading day. In comparison, we obtain a substantial improvement in precision by averaging across nearly 1000 trading days.

IV. Testing for Intraday Trading Invariance

This section presents tests for the invariance relationships and alternative specifications of the interdependencies of the market activity variables exploiting the high-frequency, data-based regression approach developed in Section II.

A. Invariance in Transaction Count by Minute-of-Day

We initially explore the suggestive relation developed in Sections II.C and II.E for alternative theories, linking the transaction count to the composite trading activity variable W . We focus on the regressions for the transaction count in equations (21), (22), and (23), which capture the interdependencies across the intraday pattern predicted by the invariance, MDH-V, and MDH-N specifications, respectively.

Figure 5 provides the intraday scatter plot of n_t versus w_t and the estimated regression line for the intraday invariance hypothesis. The one-minute observations are averaged across all days in the sample. The model (21) predicts a slope of $2/3$, represented by the dashed line. Indeed, the dashed line is visually indistinguishable from the fitted regression line in Figure 5, indicating an excellent fit. The observations corresponding to the Asian, European, and American segments in Figure 5 are depicted in blue, green and red, respectively. It is evident that the pattern for each regime is overall consistent with the invariance predictions.

The quality of the fit in Figure 5 is so good that it invites scrutiny of the few observations that seem to deviate, however slightly, from the theoretical prediction. In particular, a cluster of 16 points, marked by red crosses, falls below the regression line. These observations stem from the period 15:00–15:15, when the cash market has just closed, but the futures market remains open. Similarly, the blue crosses signifying observations from the start of daily trading, 17:01–17:06, also fall below the line. Remarkably, the slope generated by these red and blue crosses is also nearly identical to $2/3$. This parallel shift suggests that the invariance principle continues to apply, but there may be a change in the composition of active traders which alters the frequency of small versus large transactions, and thus the intercept.

For the alternative models (22) and (23), we also expect a regression slope of $2/3$. The points in Figure 6, representing the regression relation for the MDH theories, align themselves nicely along straight lines with an associated $R^2 \gg 0.99$. The fitted slopes, however, differ substantially from $2/3$. Hence, the closeness of the actual slope to the predicted one provides qualitative support to the intraday invariance principle relative to the other hypotheses.

Although the two MDH regressions generate slopes that are inconsistent with the underlying MDH theory, the strong linear relationships in Figure 6 are striking. It turns out that this feature is a direct consequence of the intraday trading invariance relation and the extreme multicollinearity among the individual (logarithmic) activity variables. Specifically, separately regressing n_t on q_t and s_t yields slope coefficients of 2.74 and 3.06, respectively, with corresponding R^2 values of 0.95 and

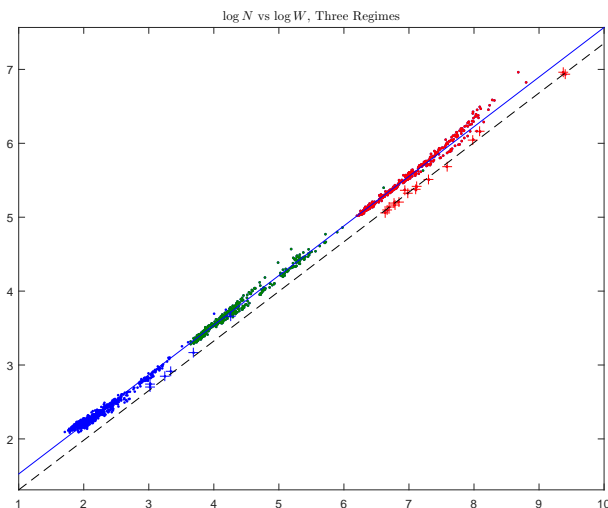


FIGURE 5. THIS FIGURE PLOTS n_t VERSUS w_t (INVARIANCE). THE BLUE CROSSES INDICATE THE FIRST 6 MINUTES OF TRADING; THE RED CROSSES INDICATE THE LAST 16 MINUTES OF TRADING. THE SOLID LINE REPRESENTS THE OLS REGRESSION; THE DASHED LINE HAS THE IDENTICAL SLOPE BUT IS FITTED TO THE RED CROSSES ONLY.




FIGURE 6. PLOTS OF n_t VERSUS $w_t - 3/2q_t$ (MDH-V) AND $w_t - q_t$ (MDH-N). THE OLS REGRESSION LINES ARE SOLID, THE MODEL PREDICTED LINES WITH SLOPE $2/3$ ARE DASHED. THE BLUE CROSSES INDICATE THE FIRST 6 MINUTES OF TRADING; THE RED CROSSES INDICATE THE LAST 16 MINUTES OF TRADING.

0.97. Thus, to a good empirical approximation, we have

$$(25) \quad n_t \approx c + 3q_t \quad \text{and} \quad s_t \approx c + q_t.$$

Heuristically, we can now demonstrate why the MDH regressions reveal an almost perfect linear dependence among the variables, yet produce slope coefficients that differ from the values implied by the underlying MDH theory. First, note that the relationships above imply that we can express w_t (approximately) in terms of q_t ,

$$(26) \quad w_t = c + n_t + q_t + \frac{1}{2}s_t \approx c + 3q_t + q_t + \frac{1}{2}q_t = c + \frac{9}{2}q_t.$$

Not surprisingly, given our prior findings, captured in Figure 5, these empirically observed linear dependency structures are consistent with intraday invariance. To see this, it is convenient to introduce some generic notation for the log-linear intraday invariance regression, labeling the slope coefficient β_I ,

$$n_t = c + \beta_I \cdot w_t.$$

Next, observe that we may approximate the left hand side via equation (25) and the right-hand side via equation (26), obtaining

$$c + 3q_t \approx c + \beta_I \left[\frac{9}{2}q_t \right].$$

It follows that $3 \approx \beta_I \cdot \frac{9}{2}$, implying $\beta_I \approx \frac{2}{3}$, which is the exact value of the slope coefficient implied by the invariance principle.

For the MDH-V hypothesis, letting β_V denote the associated regression slope coefficient, we find that equation (22) and the log-linear empirical approximations imply

$$c + 3q_t \approx c + \beta_V \left[w_t - \frac{3}{2}q_t \right] \approx c + \beta_V \left[\frac{9}{2}q_t - \frac{3}{2}q_t \right] = c + \beta_V \cdot 3q_t.$$

Hence, $\beta_V \approx 1$ follows from the log-linear approximations and the invariance relation.

Similarly, for the MDH-N hypothesis, letting β_N denote the associated regression coefficient in the specification (23), we obtain

$$c + 3q_t \approx c + \beta_N [w_t - q_t] \approx c + \beta_N \left[\frac{9}{2}q_t - q_t \right] = c + \beta_N \cdot \frac{7}{2}q_t,$$

implying $3 \approx \beta_N \cdot \frac{7}{2}$, which is equivalent to $\beta_N \approx \frac{6}{7}$.

Table 2 reports the point estimates $\beta_V = 0.98$, $\beta_N = 0.85$, and $\beta_I = 0.67$, consistent with the implications of intraday trading invariance $\beta_V \approx 1$, $\beta_N \approx \frac{6}{7} \approx 0.86$, and $\beta_I \approx \frac{2}{3}$.

TABLE 2—INTRADAY REGRESSION OF $\log N$ ON (ADJUSTED) $\log W$ FOR ALTERNATIVE MODELS

	Nobs	c	β	\bar{R}^2	β_{model}	β_{inv}
MDH-V	1335	2.41	0.98	0.997	2/3	1
MDH-N	1335	1.75	0.85	0.999	2/3	6/7
Invariance	1335	0.85	0.67	0.998	2/3	2/3

This table concerns intraday OLS regressions of n_t onto $w_t - 3/2q_t$ (MDH-V), $w_t - q_t$ (MDH-N), and w_t (Invariance). Nobs is the number of intraday observations. In the last two columns, β_{model} is the slope coefficient predicted by each model, while β_{inv} is the coefficient predicted by invariance and empirical approximation in (25).

In spite of our inability to conduct formal inference on the regressions considered in this subsection, the evidence strongly favors the intraday invariance hypothesis. The strong correspondence between theory and empirical results suggests that the

slope coefficients are subject only to minor biases, reflecting the substantial error reduction achieved through our averaging procedures. This is corroborated by the extreme R^2 values attained by each individual regression. Even if the degree of explained variation is artificially boosted by this regression set-up, large idiosyncratic errors in q_t and s_t would lower the R^2 . Furthermore, the strong rejections of the MDH specifications demonstrate that the framework readily discriminates between the alternative hypothesis.

B. Invariance in Transaction Volatility by Minute-of-Day

The encompassing regression (24) offers a more traditional setting for hypothesis testing in which standard inference techniques are applicable. The specification also highlights the dramatically different predictions regarding the relation between volatility per transaction and trade size, ranging from a positive association (MDH-V, $\beta = 1$), no relation (MDH-N, $\beta = 0$), to a strongly negative association (invariance, $\beta = -2$). This renders regression (24) a particularly clean test of the relative explanatory power of the different hypotheses.

We estimate the regression lines across the intraday pattern, with the log values of the relevant variables for each one-minute observation averaged across all days in the sample. The results are presented in Figure 7 and Table 3.




FIGURE 7. LEFT PANEL: INTRADAY SCATTER PLOT OF $s_t - n_t$ VERSUS q_t . THE OLS REGRESSION LINE IS SOLID AND THE PREDICTED INVARIANCE LINE IS DASHED. THE BLUE CROSSES INDICATE THE FIRST 6 MINUTES OF TRADING AND THE RED CROSSES THE LAST 16 MINUTES OF TRADING. RIGHT PANEL: THE SAME SCATTER PLOT AS IN THE LEFT PANEL, EXCEPT THE MINUTES AROUND THE REGIONAL MARKET OPENINGS AND CLOSURES ARE REMOVED.

The left panel of Figure 7 reveals that the 1,335 separate one-minute observations are less tightly centered on the regression line than in Figure 6. This is expected, as we no longer absorb part of the variation of the regressand in the right-hand-side variable. Nonetheless, the relation remains strikingly close to log-linear, with an adjusted R^2 of 0.966. Moreover, the estimated regression line associated with the intraday invariance hypothesis (solid, slope -2.0045) is almost identical to the theoretical line (dashed, slope -2). Meanwhile, the MDH-V and MDH-N models are overwhelmingly rejected: the regression line is neither flat nor upward sloping.

TABLE 3—INTRADAY REGRESSION OF $\log \frac{\sigma^2}{N}$ ONTO $\log Q$

	Nobs	c	β	se(c)	se(β)	\bar{R}^2
Unfiltered	1335	-2.61	-2.00	0.021	0.010	0.966
Filtered	1273	-2.59	-2.01	0.016	0.008	0.980

This table reports on the Intraday OLS regression, $s_t - n_t = c + \beta \cdot q_t + u_t$. MDH-V, MDH-N and invariance predict $\beta = 1$, $\beta = 0$, and $\beta = -2$, respectively.

We further note that most of the outliers in the left panel stem from periods during which active trading is transitioning from one trading regime to another. If we eliminate observations corresponding to 3 minutes around 1:00 and 2:00 (the beginning of European hours), the 3 + 30 minutes around 8:30 (the beginning of American hours and the 9:00 news announcements) as well as the period 15:00-15:15 (the end of American hours, following the closure of the cash market), then we obtain the results depicted in the right panel of Figure 7. The intraday invariance predictions are again validated in the sense that the regression slope equals the predicted -2 . At the same time, the vast majority of the outliers are eliminated by this filtering procedure, suggesting that intraday trading invariance offers a particularly compelling account of the interaction among the activity variables in a stable market setting.

There are a variety of ways to assess the robustness of the findings by focusing on specific subsamples. The main complication is the presence of non-trivial measurement errors in our high-frequency observations. In particular, if the regression design reduces the genuine variation of the regressor relative to the size of the measurement errors, then we will mechanically find the regressions to have less explanatory power. Moreover, we may start observing a distinct downward bias in the (absolute) size of the estimated regression slope due to the errors-in-variables effect. With these caveats in mind, we first investigate the performance of the invariance hypothesis over annual subsamples. Due to the stability of the diurnal pattern for the activity variables, this approach retains the pronounced variation in the regressor, so the primary effect is simply a loss of observations.

Panel A of Table 4 reports the results for the year-by-year log-linear regressions. The individual regression slopes are now scattered more widely, but there is no evidence of any systematic deviation from the theoretical value of -2 implied by the intraday invariance hypothesis.

An alternative robustness check involves regime-by-regime regressions. This approach does reduce the sample variation of the regressor quite substantially, raising concerns about the impact of measurement error. Panel B of Table 4 reports the regression results for all observations within the individual regimes, while Panel C reports the results when some observations associated with the transition periods between regimes are excluded, as in Table 3.

In Panel B, we now observe a clear drop in the (negative) slope coefficients. Panel C shows that filtering improves matters considerably, but still leaves a non-trivial

TABLE 4—TIME-SERIES REGRESSION OF $\log \frac{\sigma^2}{N}$ ONTO $\log Q$

Panel A: 4 Years						
	Nobs	c	β	$se(c)$	$se(\beta)$	\bar{R}^2
2008	1335	-2.89	-1.88	0.021	0.011	0.958
2009	1335	-2.23	-2.13	0.023	0.012	0.961
2010	1335	-2.27	-2.17	0.032	0.016	0.936
2011	1335	-3.28	-1.74	0.026	0.012	0.939
All	1335	-2.61	-2.00	0.021	0.010	0.966

Panel B: 3 Regimes, Unfiltered						
	Nobs	c	β	$se(c)$	$se(\beta)$	\bar{R}^2
Asia	540	-3.90	-1.14	0.053	0.034	0.672
Europe	390	-3.16	-1.73	0.092	0.046	0.784
America	405	-5.03	-1.06	0.176	0.070	0.361
Pooled	1335	-2.61	-2.00	0.021	0.010	0.966

Panel C: 3 Regimes, Filtered						
	Nobs	c	β	$se(c)$	$se(\beta)$	\bar{R}^2
Asia	526	-3.84	-1.18	0.051	0.034	0.704
Europe	388	-3.10	-1.76	0.086	0.043	0.810
America	359	-3.21	-1.78	0.182	0.073	0.623
Pooled	1273	-2.59	-2.01	0.016	0.008	0.980

This table reports on OLS regressions of the form $s_t - n_t = c + \beta \cdot q_t + u_t$. In Panel A, the coefficients, standard errors, and \bar{R}^2 statistics are estimated separately for each calendar year and then for the whole sample. The last year ends November 4, 2011. In Panels B and C, the regressions are estimated separately for each regime, except that in Panel C the minutes around the regional market openings and closures are removed. MDH-V, MDH-N and invariance predict $\beta = 1$, $\beta = 0$, and $\beta = -2$, respectively.

gap between the estimated slopes and -2 . The discrepancy is especially striking for the Asian regime, where the diurnal pattern, and thus the systematic intraday variation of the regressor, is less pronounced compared to the other regions, and market activity is generally somewhat subdued.

As a final robustness check, we test the invariance relation within each quarter in our sample. Figure 8 displays the slope coefficient for the q_t variable across the quarters. It is evident that the slope coefficient is economically close to the theoretical value of -2 in all instances.¹⁵

In summary, our findings are broadly consistent with the intraday invariance principle, while the MDH specifications are decidedly rejected. In particular, there is strong support for the hypothesis that market participants actively adjust their

¹⁵The confidence bands are obtained assuming we only face sampling error, but our construction of expectation proxies through the MEM approach inevitably introduces non-trivial measurement errors into the specification. While it is difficult to quantify this effect, it is clear that the confidence bands are too narrow and should be viewed with skepticism.

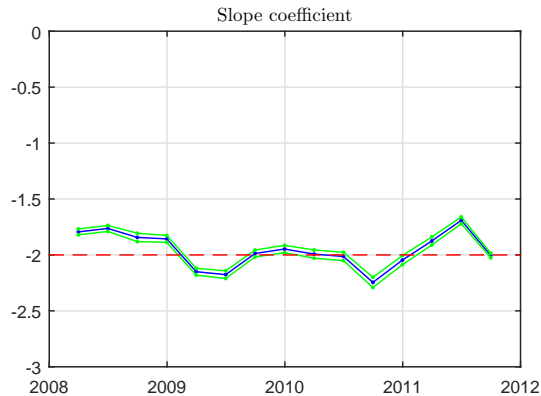


FIGURE 8. THE FIGURE PLOTS THE SLOPE COEFFICIENT β FOR OLS REGRESSIONS $s_t - n_t = c + \beta \cdot q_t + u_t$ ESTIMATED SEPARATELY FOR EACH QUARTER. ALSO SHOWN ARE CONFIDENCE INTERVALS COMPUTED AS ± 2 STANDARD ERRORS. THE RED DASHED LINE INDICATES THE THEORETICAL VALUE PREDICTED BY INVARIANCE.

trade size in response to varying market conditions in a manner consistent with the intraday invariance specification. Nonetheless, it is also evident that temporary deviations from the predicted relationships occur when the trading environment is changing rapidly. A detailed exploration of the breakdown in invariance during such market transitions falls outside the scope of the current paper, but we do present some additional evidence on this issue in Section V.C.

C. Invariance in Daily Time-Series Observations

The log-linear regressions in Section IV.B exploit the systematic minute-by-minute variation in the activity variables across the diurnal pattern to generate a powerful test of the intraday-invariance hypothesis. As an empirical matter, we also observe large and persistent fluctuations across days, as illustrated in Figure 2, albeit the variation is substantially smaller than observed intraday. In fact, most existing studies of the interaction between transaction activity and return volatility rely on daily data, rendering tests based on the intraday pattern infeasible.

In this section, we contrast the performance of the intraday-invariance hypothesis to the MDH specifications in the time-series dimension. We achieve this by aggregating the one-minute estimates for the activity variables for each day and each trading regimes to obtain day-regime values, following the scheme in equation (18).

In the preceding section, our specification of intraday invariance left out the price variation. When averaging across days to examine minute-by-minute variation within days, this makes no difference. When averaging across minutes within days to examine the time-series variation across days, the effect may be significant, however, because the level of the S&P 500 index varies greatly over the sample period. Given the low level of inflation, nominal fluctuations in the S&P 500 index represent fluctuations in the real value of assets in the economy and may also represent fluctuations in the real income and productivity of finance professionals.

These effects may partially offset each other if the value of the futures contract covaries with the real wealth of market participants. We retain our prior specification of the intraday invariance hypothesis below. In Appendix VI.B, we present results for alternative representations that include the asset price in the regression. We find such formulations to under-perform the specifications analyzed below uniformly. A more detailed study of these effects is an interesting issue for future research.

The daily time series for each of the three trading regimes are distinct, with each clearly being affected by its specific regional patterns of trading activity and economic news releases. In other words, we effectively have three alternative daily series available for our tests. Thus, for each of the intraday activity variables, we generate three (day-regime) data points per trading day by averaging separately each day across all intraday observations within the sets, $\mathcal{D} \times \mathcal{T}_1$, $\mathcal{D} \times \mathcal{T}_2$, and $\mathcal{D} \times \mathcal{T}_3$, respectively. The day-regime series (averaged across a singular regime) are likely to be noisier than the corresponding intraday series (averaged across all trading days) explored in Section IV.A because each regime contains fewer observations than there are trading days in our sample. Furthermore, the time-series variation is less pronounced than the intraday variation. Hence, compared to the intraday tests, the genuine variation in the regressor—the signal strength—is smaller relative to the measurement errors in the time-series setting. As a consequence, we expect a larger degree of dispersion around the regression line and we may note a downward bias in the regression slope due to the classical errors-in-variable effect. To gauge the severity of these distortionary features, we also perform a joint test, where we enhance the variation of the regressor by including all three day-regime time series in one pooled regression.

Figure 9 presents the results for a regression of log variance per transaction on log trade size, constructed along the lines of equation (24). It may be compared to the evidence in Figure 7. Figure 9 shows that the points cluster along a fitted line with slope -1.99 , close to the predicted slope of -2 , for the regression pooling all the day-regime observations. It is also evident that the dispersion around the regression line increases relative to the intraday regression displayed in Figure 7.

As we move to the regime-specific regressions, the noise-to-signal ratio rises, and Panel A of Table 5 documents that the regression slopes decline in absolute size, consistent with an error-in-variables bias and relatively large heterogeneity across the observations within regimes. Nonetheless, the evidence remains strongly in favor of the intraday invariance principle relative the MDH hypotheses which imply either a flat or positively-sloped regression line. This conclusion is corroborated by Panel B of Table 5 for annual regressions based on the pooled day-regime observations. The slope coefficients for the individual years are all close to -2 , except for the slightly lower (absolute) value obtained over the abbreviated final year of the sample, 2011. This is consistent with an enhanced signal-to-noise ratio, obtained by incorporating the variation in the activity variables across regimes, enabling a more precise inference, even for moderately sized samples.

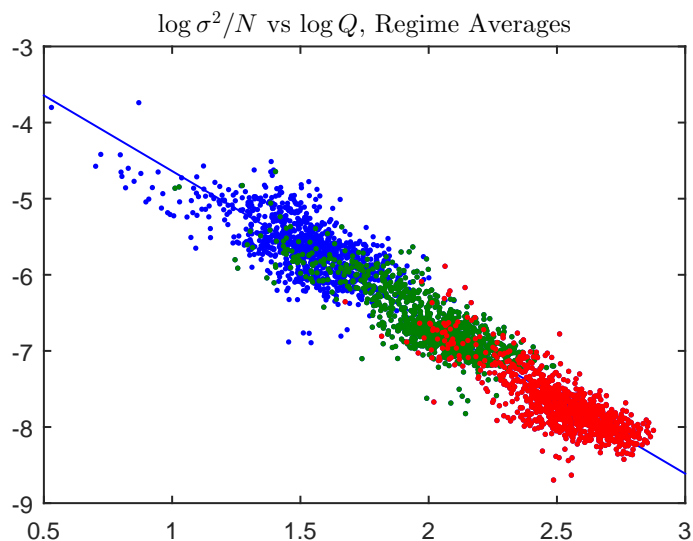


FIGURE 9. SCATTER PLOT OF $s_{d,r} - n_{d,r}$ VERSUS $q_{d,r}$. FOR EACH DAY d , THERE ARE THREE OBSERVATIONS CORRESPONDING TO THE THREE REGIMES $r = 1, 2, 3$. FOR EACH VARIABLE, WE FIRST COMPUTE THE LOGARITHMIC VALUE AND THEN AVERAGE IT ACROSS EACH REGIME FOR A GIVEN DAY.

TABLE 5—TIME-SERIES REGRESSION OF $\log \frac{\sigma^2}{N}$ ONTO $\log Q$

Panel A: 3 Regimes						
	Nobs	c	β	$se(c)$	$se(\beta)$	\bar{R}^2
Asia	965	-3.54	-1.38	0.076	0.050	0.445
Europe	969	-3.22	-1.70	0.070	0.035	0.708
America	969	-3.32	-1.74	0.101	0.040	0.661
Pooled	2903	-2.65	-1.99	0.023	0.011	0.918
Panel B: 4 Years						
	Nobs	c	β	$se(c)$	$se(\beta)$	\bar{R}^2
2008	748	-2.86	-1.90	0.035	0.018	0.938
2009	754	-2.28	-2.11	0.039	0.020	0.938
2010	756	-2.38	-2.12	0.047	0.022	0.923
2011	645	-3.29	-1.73	0.049	0.022	0.902
All	2903	-2.65	-1.99	0.023	0.011	0.918

This table reports on the OLS regressions $s_{d,r} - n_{d,r} = c + \beta \cdot q_{d,r} + u_{d,r}$. For each day, there are three observations corresponding to the regimes $r = 1, 2, 3$. In Panel A, the coefficients, standard errors, and \bar{R}^2 statistics are estimated separately for each regime. In Panel B, the regressions are estimated separately for each calendar year. The last year ends November 4, 2011. MDH-V, MDH-N and invariance predict $\beta = 1$, $\beta = 0$, and $\beta = -2$, respectively.

D. Invariance for Multi-Minute Bin Aggregation

Our empirical tests for intraday trading invariance depend critically on high-frequency measurement of trading intensity and return volatility. The minute-

by-minute measures are invariably noisy. So far, we have performed aggregation across trading days to generate fairly accurate observations of the systematic intraday pattern, and we have aggregated across the intraday observations to obtain time series comprised of reasonably precise daily measures. Either approach involves a compromise where we sacrifice some systematic variation in the data along specific dimensions in order to improve the precision.

In this section, we seek to retain more of the high-frequency variation in the activity variables, while accepting that the individual measures become noisier. Specifically, we aggregate the one-minute observations into 31 separate bins. Each bin contains 90 consecutive minutes during the Asian regime, 39 minutes in the European regime, and 27 minutes in the American regime. This binning approach allows the measures to reflect both the intraday variation in trading activity across time bins and the time-series variation across days. The procedure generates 6, 10 and 15 day-bin observations for each day across the three regimes, respectively. Furthermore, it ensures that the regimes are relatively homogeneous in terms of the magnitude of the sampling and measurement errors. Following the scheme in equation (19), the subscript b refer to a bin within the trading day so, for example, $n_{d,b}$ denotes the log number of trades per minute within the bins. In total, we have nearly 30,000 day-bin ($D \times 31$) observations for the relevant activity variables, i.e., 31 observations for each of the variables over $D = 969$ days.¹⁶

Figure 10 depicts the scatter plot of $s_{d,b} - n_{d,b}$ versus $q_{d,b}$ for all day-bin observations along with the regression line obtained from the joint fit to all the data points. Consistent with the presence of non-trivial measurement errors, the points are now much more widely scattered around the regression line compared with Figure 9. Nonetheless, as reported in Table 6, the slope estimate equals -1.84, which is close to the theoretical value of -2 implied by invariance. Likewise, for the individual years, the corresponding regression slopes attain a similar magnitude. In other words, the large degree of variation in the regressor across the bins may be sufficient to alleviate the severity of the errors-in-variable bias, both for the joint and the year-by-year regressions.

V. Invariance during Market Transitions

Certain events corresponding to particularly turbulent market conditions generate large outliers. In this section, we explore whether the intraday invariance principle can accommodate these dramatic episodes of elevated market activity.

A. Macroeconomic Announcements

We now examine whether intraday invariance provides a good characterization of the market dynamics surrounding releases of macroeconomic announcements.

¹⁶A few observations are lost due to market closures or failures in the dissemination of data.




FIGURE 10. SCATTER PLOT OF $s_{d,b} - n_{d,b}$ VERSUS $q_{d,b}$ WITH DATA AGGREGATED ACROSS THE ONE-MINUTE OBSERVATIONS WITHIN BIN b FOR EACH TRADING DAY d . AS BEFORE, THE THREE REGIMES ARE REPRESENTED BY DISTINCT COLORS.

TABLE 6—TIME-SERIES REGRESSION OF $\log \frac{\sigma^2}{N}$ ONTO $\log Q$ USING BINNED DATA

Panel A: 3 Regimes						
	Nobs	c	β	$se(c)$	$se(\beta)$	\bar{R}^2
Asia	5790	-4.05	-1.03	0.027	0.017	0.377
Europe	9666	-3.78	-1.42	0.026	0.013	0.559
America	14499	-4.29	-1.36	0.032	0.013	0.439
Panel B: 4 Years						
	Nobs	c	β	$se(c)$	$se(\beta)$	\bar{R}^2
2008	7705	-3.16	-1.76	0.017	0.008	0.859
2009	7779	-2.60	-1.96	0.020	0.009	0.846
2010	7806	-2.93	-1.88	0.025	0.011	0.788
2011	6665	-3.51	-1.64	0.025	0.011	0.783
All	29955	-2.98	-1.84	0.011	0.005	0.822

The table reports on OLS regressions $s_{d,b} - n_{d,b} = c + \beta \cdot q_{d,b} + u_{d,b}$ across days d and bins b . In Panel A, the coefficients, standard errors, and \bar{R}^2 statistics are estimated separately for each regime. In Panel B, the regressions are estimated separately for each calendar year. The last year ends on November 4, 2011. MDH-V, MDH-N and invariance predict $\beta = 1$, $\beta = 0$, and $\beta = -2$, respectively.

For the U.S. market, the most important announcements occur at 7:30. The Employment Report is usually released at 7:30 on the first Friday of the month and the Consumer Price Index at 7:30 on the second Friday of the month. Other




FIGURE 11. THE FIGURE DEPICTS AVERAGES ACROSS TRADING DAYS FOR WHICH AN ASSUMED 7:30 CT ANNOUNCEMENT INDUCED THE LARGEST SUBSEQUENT INCREASE IN TRADING ACTIVITY. THE STATISTICS INCLUDE CONTRACT VOLUME V_t (PER MINUTE), VOLATILITY σ_t , AVERAGE TRADE SIZE Q_t , AND THE NUMBER OF TRADES. THE AVERAGES ARE COMPUTED AT THE GRANULARITY OF ONE MINUTE. THE DASHED VERTICAL LINES SEPARATE THE TRADING REGIMES. THE SOLID VERTICAL LINE INDICATES THE TIMING OF THE ANNOUNCEMENT.

releases at this time of day include the Producer Price Index, the Employment Cost Index, the U.S. Import/Export Price Indices, and Real Earnings. Market activity variables often exhibit distinct spikes immediately after such releases. Rationalizing the interaction among the trading variables and return volatility in the periods surrounding these events constitutes a challenging test for intraday invariance.

To identify the relevant announcements, we focus on the days with the largest increase in trading activity immediately after 7:30. Specifically, we compute the ratio of the number of trades for the 10 minutes after 7:30 to the number of trades for the 60 minutes before 7:30. We then select the ten percent of days (97 out of 969) with the highest ratio. On average, this procedure identifies about two days each month with significant spikes in trading activity, likely associated with 7:30 announcements. In line with the terminology of Section II.D, we denote this subset of days by \mathcal{D}_A . We also zoom in further on the announcements and study the three minutes before and after 7:30; we denote this subset of minutes by \mathcal{T}_A .

Figure 11 depicts the intraday statistics averaged across the announcement days,

\mathcal{D}_A . Compared to the full sample statistics in Figure 1, the 7:30 spike is much more pronounced for volume, number of trades, and volatility. Close inspection also reveals a small, yet distinctive, downward shift in the average trade size.




FIGURE 12. THE FIGURE PROVIDES SCATTER PLOTS FOR $(s_t - n_t)$ VERSUS q_t . THE LEFT AND RIGHT PANELS PLOT AVERAGES ACROSS ALL DAYS AND 7:30 ANNOUNCEMENT DAYS, RESPECTIVELY. THE 3 MINUTES BEFORE 7:30 ARE INDICATED BY RED DOTS, THE 3 MINUTES THEREAFTER BY A BLUE ASTERISK AND TWO BLUE CROSSES, AND ALL OTHER MINUTES BY LIGHT GRAY DOTS. THE REGRESSION LINE FROM FIGURE 7 IS ALSO SHOWN.

Figure 12 displays the regression fit for the intraday invariance relation, as captured by the plot of the averaged one-minute observations for volatility per transaction versus trade size. The left panel includes all trading days, while the right panel contains only the 7:30 announcement days. In both panels, the three minutes preceding 7:30 are represented by red dots, the 7:30-7:31 announcement observation by the blue asterisk, and the following two one-minute intervals by the blue crosses. The left panel reveals that none of the points surrounding the announcement generally represent outliers. On announcement days, the observations around the release time remain in line with theoretical predictions, except for the exact interval covering the news release (blue asterisk). The latter observation lies above the regression line, indicating excess volatility relative to trading volume. This may happen if the price jumps at the release without any trading, consistent with the observed market dynamics around scheduled releases. The bid-ask spreads widen and trading stalls in the moments before the release. After the number is reported, the quotes jump discretely in response to the news content, and only thereafter does an avalanche of trading hit the market. In short, the public release of news at predetermined times allows the headline number to be incorporated into quotes prior to any significant trading. Only subsequently is the price effectively driven by the order flow. Modeling trading during announcements may require a careful adjustment for the fraction of return volatility unrelated to order flow.

Figure 13 present the corresponding plots for the 9:00 announcement days. The releases at 9:00 include New and Existing Home Sales, the Housing Market Index, Consumer Sentiment, and Business Inventories, among others. On average, these




FIGURE 13. THE FIGURE PROVIDES SCATTER PLOTS FOR $(s_t - n_t)$ VERSUS q_t . THE LEFT AND RIGHT PANELS PLOT AVERAGES ACROSS ALL DAYS AND 9:00 ANNOUNCEMENT DAYS, RESPECTIVELY. THE 3 MINUTES BEFORE 9:00 ARE INDICATED BY RED DOTS, THE 3 MINUTES THEREAFTER BY A BLUE ASTERISK AND TWO BLUE CROSSES, AND ALL OTHER MINUTES BY LIGHT GRAY DOTS. THE REGRESSION LINE FROM FIGURE 7 IS ALSO SHOWN.

announcements have a less pronounced effect on average trade size.¹⁷ Nevertheless, the spikes in the average for volume, number of trades, and volatility at the announcement time remain dramatic. As before, the figure confirms that intraday invariance holds up well, even during such unusual market conditions. Since the number of announcement days is limited and the measurement errors are likely larger around the release time, the small downside deviation for the two blue crosses is immaterial. As before, the observation at the exact release interval (blue asterisk) lies above the subsequent blue crosses. This may again reflect the tendency for prices to jump in the absence of trading at the moment the announcement hits the news wires.

We conclude that, broadly speaking, the intraday invariance principle provides an accurate description of the interdependencies among the market activity variables, even around news releases, when the fluctuations are magnified greatly.

B. The Flash Crash

We next demonstrate that intraday invariance provides a framework that may be helpful in interpreting the chaotic market conditions during the flash crash of May 6, 2010. The study of this market crash may help to understand the underlying dynamics behind other episodes of extreme market behavior. Kyle and Obizhaeva (2016a) describe the flash crash as well as other stock market crashes associated with execution of large bets. Andersen and Bondarenko (2014, 2015) discuss other early warning signals for market turbulence.

The top left panel of Figure 15 depicts the price dynamics of the E-mini S&P 500 futures on May 6, 2010. The vertical lines indicate the timing of the crash from

¹⁷The intraday variation of the activity variables around the 9:00 news releases are qualitatively similar to those observed for the 7:30 releases in Figure 11, so we omit this figure.

13:32 to 13:45 CT, as identified in the joint report by Staffs of the CFTC and SEC (2010a,b). In the morning, the market declined by about 3 percent. There were rumors about a debt default by Greece, elections in the U.K., and an upcoming jobs report in the U.S. From 13:40 to 13:45, prices first plummeted by 5.12% and then recovered by 5% over the next ten minutes, after a pre-programmed circuit breaker within the CME Globex electronic trading platform halted trading for five seconds. Prices entered a free fall only during the last minute of the event window. The crash was accompanied by record trading volume and extreme realized volatility, as seen in the top right and bottom left panels of Figure 15. The apparent breakdown in the provision of market liquidity is obviously of great interest.

We next examine whether intraday invariance relationships held during the market crash or the data displayed some unusual patterns. Since the crash event lasted for less than an hour, we focus on the prediction that the invariant $\log I_{dt}$ is an identically and independently distributed random variable over this period, as suggested by equation (8). We test this proposition informally by calculating the time series of average $\log I_{dt}$ computed at a four-minute frequency. We further normalized those variables to have a zero mean and a unit variance, using observations of the trading invariant across all days at the same point in time.¹⁸ We then simply look for outliers in the intraday sequence on May 6, 2010.

The dynamics of the standardized unsigned trading invariant $\log I_{dt}$ is shown in the lower right panel of Figure 15. There are three consecutive large outliers—exceeding 3 standard deviations each—during the minute before and the first eleven minutes immediately following the five-second trading halt at the bottom of the flash crash. When compared with realizations of $\log I_{dt}$ during the same time interval for other days in the sample, these three outliers represent 100th, 99.9th, and 99.4th percentile events. Moreover, for the four-minute interval ranging from five to one minutes before prices bottomed out, the value of the invariant $\log I_{dt}$ corresponds to a 78th percentile event. Prior to these outliers, there were no observations of $\log I_{dt}$ exceeding ± 2 standard deviations. In summary, the fluctuations in $\log I_{dt}$ were contained within the typical range both prior to the crash and during the time when prices dived, but the invariant attained extreme positive values at the exact time when the market collapsed and then, after the five-second market break, as prices started to recover.

How can we relate the apparent breakdown in intraday invariance, suggested by these consecutive outliers, to any hypothesis regarding the causes for the collapse in liquidity provision during the flash crash?

Menkveld and Yueshen (2015) document that the co-integrating relation between the E-mini contract and the SPDR S&P 500 exchange traded fund (the “spider” or SPY) broke down about one minute prior to the trading halt and resumed approximately eight minutes afterwards. The three outliers in $\log I_{dt}$ coincide exactly with

¹⁸The four-minute blocks represent a compromise between the desire for short blocks to capture the rapidly changing market dynamics and long blocks to facilitate averaging out measurement errors.




FIGURE 14. THE FIGURE SHOWS PRICE P , VOLUME V , VOLATILITY σ , AND TRADE SIZE Q ON MAY 6, 2010. THE STATISTICS ARE COMPUTED AT A GRANULARITY OF ONE MINUTE. THE SOLID VERTICAL LINES INDICATE THE TIMING OF THE FLASH CRASH.




FIGURE 15. FOR MAY 6, 2010, THIS FIGURE DEPICTS THE LOG INVARIANT (LEFT PANEL) AND THE LOG INVARIANT STANDARDIZED TO HAVE ZERO MEAN AND UNIT VARIANCE (RIGHT PANEL) AT A 4-MINUTE FREQUENCY. THE RED LINES IN THE LEFT PANEL REPRESENT ± 2 STANDARD DEVIATION BOUNDS COMPUTED USING OBSERVATIONS ACROSS ALL DAYS. THE SOLID VERTICAL LINES INDICATE THE TIMING OF THE FLASH CRASH.

this drop in liquidity provision. Several factors hampered the implementation of cross-market arbitrage. Margin requirements rose rapidly in response to the spike

in volatility, there were problems with connectivity across trading platforms, and there was a general sense of confusion among market participants.

Kirilenko et al. (2014) report (see their Table II) that the average trade size of groups of traders which represent institutional investors (Fundamental Buyers and Fundamental Sellers) is larger than the average trade size of the groups that intermediate the trading process (High-Frequency Traders, Market Makers, and Opportunistic Traders). In the last few minutes prior to the nadir of the flash crash, the participation rate of high-frequency traders increased, but then dropped immediately as prices bottomed out (see their Figure 6).

Taken together, these results suggest that the breakdown in arbitrage relationships, which occurred right around the time prices hit the bottom during the flash crash, is associated with a low participation rate by market makers and arbitrageurs. Since these intermediaries tend to execute orders of small size, we expect the average trade size to rise when their participation rate falls, leading to larger values for the trading invariant $\log I_{dt}$. We therefore conjecture that the apparent breakdown in intraday trading invariance relationships, captured through the three large outliers in $\log I_{dt}$, may indicate a lower participation rate of market makers and arbitrageurs than during normal market conditions. It does not imply a breakdown in the invariance of bets, as discussed by Kyle and Obizhaeva (2016c) and Kyle and Obizhaeva (2016a).

C. The Cash Market Open and Close

Under the null hypothesis of intraday trading invariance, we can infer the implied average trade size as a function of the expected volume and return volatility, both of which may be estimated by the techniques developed in Section II.C. Comparing the predicted trade sizes with actual minute-by-minute observations of average trade size, which may exhibit highly idiosyncratic features at specific times during the daily cycle, represents another challenging test for intraday invariance.

From the basic intraday invariance relation (24) with $\beta = -2$ and $v = n + q$, we obtain the theory-implied trade size q_t^* given by,

$$(27) \quad q_t^* = c + \frac{1}{3} [v_t - s_t].$$

We identify the constant c via the (moment) condition that the average implied log trade size q_t^* matches the sample average of the realized log trade size q_t .

Figure 16 compares the actual average log trade size q_t and the corresponding implied log trade size q_t^* . The left panel shows that implied trade size generally tracks the actual trade size closely, confirming that intraday trading invariance provides a very accurate account of how the average trade size responds to shifts in volume and return volatility on a minute-by-minute level.

A set of features of the figure are noteworthy. First, we detect a distinct U-shaped pattern for average log trade size within each trading regime. The most striking deviation from the U-shaped pattern occurs from 1:00 to 2:00. This period




FIGURE 16. THE LEFT PANEL PLOTS THE EMPIRICAL TRADE SIZE q_t AND THE IMPLIED TRADE SIZE q_t^* FOR EACH MINUTE t , AVERAGED ACROSS ALL TRADING DAYS IN THE SAMPLE. THE RIGHT PANEL DEPICTS THE PREDICTION ERROR $q_t - q_t^*$. THE DASHED VERTICAL LINES SEPARATE THE THREE TRADING REGIMES.

seems to blur the boundary between the trading hours for the Asian and European cash markets. This reflects a couple of institutional features. There is no distinct opening time for European trading hours, as some exchanges in continental Europe open at 1:00, while the trading venues in London open one hour later at 2:00. Furthermore, the trading venues in Europe and America adhere to a summer time convention, whereas the major trading venues in Asia do not. This makes the boundary between Asian and European trading hours shift back and forth depending on the time of the year.

Second, in spite of the dramatic spikes in the trading volume and return volatility around the release of macroeconomic announcements, there are no apparent outliers at the typical times for those releases, such as at 7:30 and 9:00. This is consistent with the evidence in Section V.A that intraday trading invariance provides an excellent approximation to the interaction among the market activity variables during announcement periods.

Third, we observe strikingly large discontinuities in trade size when the market activity shifts from one region to another, due to the opening and closing of exchange trading in the regional cash markets. Trade size increases significantly when the European cash markets open, and it increases even more as trading shifts from Europe to America as the cash markets in that region open up.

A more detailed analysis of the deviation between the model-implied and actual trade size is difficult on the basis of the figure, since the two curves largely overlap, obscuring the gap. Hence, for visual purposes, the right panel of Figure 16 depicts the log prediction error $q_t - q_t^*$. Most of the errors fall in the range ± 0.1 , or $\pm 10\%$. Since the average trade size is less than 10 contracts, as documented in Table 1, the average prediction errors typically correspond to less than a single contract.

The right panel provides a more transparent picture of the striking discontinuities in trade size when market activity shifts from one region to another. The downward spikes in the right panel of Figure 16 at exactly 1:00, 2:00 and 8:30 represent clear

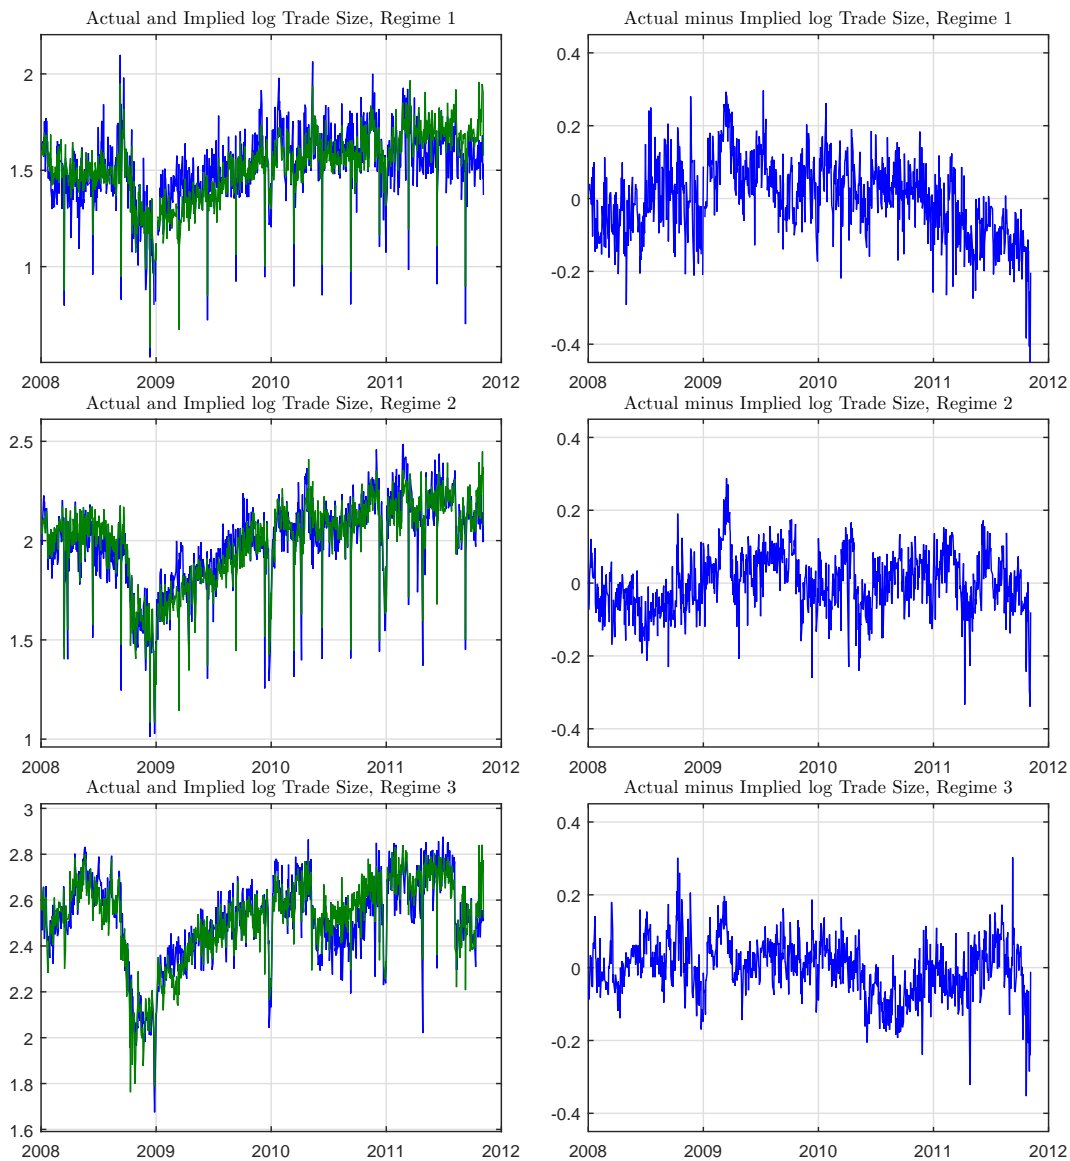


FIGURE 17. THE LEFT PANELS PLOTS THE EMPIRICAL TRADE SIZE $q_{d,r}$ AND THE IMPLIED TRADE SIZE $q_{d,r}^*$, AVERAGED ACROSS THE ONE-MINUTE OBSERVATIONS IN EACH REGIME $r = 1, 2, 3$ AND FOR EACH DAY d . THE RIGHT PANELS DEPICT THE CORRESPONDING PREDICTION ERRORS, $q_{d,r} - q_{d,r}^*$.

outliers, revealing that the actual trade sizes are somewhat lower than the implied trade size at exactly the time when a more active cash market opens. Notice that intraday invariance, correctly, predicts a sharp increase in the trade size at these points but, systematically, the actual trade size jumps somewhat less than predicted. Likewise, intraday invariance predicts a sharp increase in trade size when trading transitions to a regime where the futures market is operating but the underlying cash market is closed. As before, the directional prediction of intraday invariance is correct but, in contrast to the case when the cash market is coming

on line, the actual trade size now jumps more than predicted. These effects are visible at 15:00, as the U.S. cash market closes, and for the market opening at -7:00, i.e., 17:00 the preceding day. These abrupt shifts in trade size pose an interesting challenge for all theories of trade size.

Nonetheless, it is apparent that the intraday-invariance principle provides a much better approximation to the dynamics of trade size than the MDH hypotheses. For example, the extreme jump in the number of transactions and volume in the minutes prior to the close of the cash market at 15:00, of the order of 300% and 200%, respectively, is accompanied only by a modest increase in return volatility; this is starkly at odds with either of the MDH specifications. In contrast, intraday invariance accommodates the majority of this effect through the simultaneous sharp increase in the expected trade size.

Overall, at a qualitative level, the intraday-invariance principle provides a satisfactory account of the intraday dynamics of trade size across the trading day. At specific times, however, a substantial gap emerges between the predicted and actual trade sizes. The common trait, characterizing negative versus positive outliers, is that the cash markets are just opening for the former while they are closed or closing for the latter. It suggests that part of the explanation stems from the different population of active traders in the market during times when the cash and futures market can be arbitrated versus when the futures market is operating on its own. This reasoning is also consistent with the conjecture in Section V.B that the market malfunctions around the flash crash were accompanied by a shift in the proportion of intermediaries and arbitrageurs in the market. In other words, the actual trade size may exceed the predicted one following the U.S. close and the market opening at 17:00 because a relatively larger fraction of trades represent institutional investors. The latter execute larger trades compared to market makers and arbitrageurs, who trade more frequently but at a smaller average size.

Finally, we explore the corresponding relation between the actual and predicted log trade sizes, given in equation (27), across time for each of the three regime. The panels in the left column of Figure 17 provide time-series plots for each regime. It is evident that the lower level of aggregation is associated with more noisy measures. Nonetheless, we still clearly capture the vast majority of the temporal variation in the trade size series over the sample. The panels in the right column verify that the errors are only weakly correlated and generally quite small. The only noteworthy systematic deviation within the liquid American market segment is in the midst of 2010 and possibly at the very end of the sample. These two episodes coincide with European sovereign debt crises. Although the evidence of systematic errors is much weaker than for the intraday pattern, the finding does suggest that active investors were more cautious in terms of the exposures they took on during these turbulent periods than predicted by the intraday invariance principle.

We leave further investigation of why systematic deviations from the implied trade size occur at the open and close of active cash markets as well as, possibly, during certain longer stressful market episodes for future research.

VI. Conclusion

The invariance properties of trading patterns in the E-mini S&P 500 are unexpected and powerful results, which raise interesting challenges for both empirical and theoretical research into market microstructure.

From an empirical viewpoint, it is interesting to ask how universal our current findings are. On the one hand, even if the results apply to this important and highly liquid futures market—often characterized as the primary location for price discovery in U.S. equities—they may still be an artifact of institutional arrangements or other unique features of this specific marketplace. Possibilities include the large tick size, large contract size, tight integration with liquid cash markets, important role of high-frequency traders, presence of specialized automated trading algorithms, or specific features of the CME Globex matching engine. On the other hand, if intraday trading invariance is a more universal phenomenon for financial markets, then it speaks to deep structural issues related to how trading in financial markets operates. For example, intraday trading invariance may be a characteristic of electronic platforms on which traders shred big bets into many small orders, *and* it may also apply to non-electronic dealer markets in which bets are often executed as single trades. As such, intraday trading invariance may provide a fruitful framework for analyzing a host of issues surrounding market organization, liquidity, functional operation, general trading motives, and strategies.

From a theoretical viewpoint, our findings pose a difficult conceptual question. Intraday trading invariance clearly reflects the spirit of the market microstructure invariance hypothesis, formulated for bet arrival rates and bet sizes. There are theoretical models such as Kyle and Obizhaeva (2016*b*) consistent with that hypothesis. Yet, trades are different from bets, traders usually shred bets into multiple small trades, optimizing their strategies to control transactions costs, as modeled in Kyle, Obizhaeva and Wang (2014, 2015).¹⁹ We are aware of no theoretical research which can explain how order shredding with restrictions on tick size and lot size might lead to intraday trading invariance.

Finally, our findings pose a challenge to the literature that views return volatility as governed by a single market activity measure such as either the number of transactions or the trading volume. In contrast, we find strong support for the intraday invariance relation which involves a simultaneous adjustment in a set of activity variables as market conditions evolve. In particular, trade size is endogenous and reacts immediately to variation in volatility and trading intensity. The pronounced variation in transaction size implies that the relation between trading intensity and volume fluctuates systematically with the market environment. Consequently, the standard MDH specifications fail spectacularly in accommodating the interaction among the full set of activity variables.

It is of interest to explore whether our findings favoring the intraday-invariance

¹⁹For earlier contributions to the literature on optimal order shredding, see, e.g., Bertsimas and Lo (1998), Almgren and Chriss (2000), and Obizhaeva and Wang (2013).

hypothesis are common across assets and time or whether they stem from the particular institutional features in this market or, alternatively, they are induced by the emergence of high-frequency and algorithmic trading in electronic markets. An empirical difficulty is the lack of a consistent recording of the transaction size for the active market participant in most trading venues due to recording conventions and the increasing fragmentation of trading across competing venues.²⁰

In summary, the success of the invariance principle at the high-frequency level for the E-mini S&P 500 futures, if confirmed more generally, raises intriguing possibilities to shed new light on the intrinsic workings of modern financial markets. At a minimum, given the central position of the E-mini futures trading for the functioning of global financial markets, the issues raised pose a host of questions for future empirical and theoretical research in market microstructure and market activity dynamics.

REFERENCES

- Ait-Sahalia, Yacine, and Jean Jacod.** 2014. *High-Frequency Financial Econometrics*. Princeton University Press.
- Alexander, Gordon, and Mark Peterson.** 2007. “An Analysis of Trade-Size Clustering and its Relation to Stealth Trading.” *Journal of Financial Economics*, 84: 435–471.
- Almgren, Robert, and Neil Chriss.** 2000. “Optimal Execution of Portfolio Transactions.” *Journal of Risk*, 3: 5–39.
- Andersen, Torben G.** 1996. “Return Volatility and Trading Volume: An Information Flow Interpretation of Stochastic Volatility.” *Journal of Finance*, 1: 169–204.
- Andersen, Torben G., and Oleg Bondarenko.** 2014. “VPIN and the Flash Crash.” *Journal of Financial Markets*, 17: 1–46.
- Andersen, Torben G., and Oleg Bondarenko.** 2015. “Assessing Measures of Order Flow Toxicity and Early Warning Signals for Market Turbulence.” *Review of Finance*, 19(1): 1–54.
- Andersen, Torben G., and Tim Bollerslev.** 1997. “Intraday Periodicity and Volatility Persistence in Financial Markets.” *Journal of Empirical Finance*, 4: 115–158.
- Andersen, Torben G., and Tim Bollerslev.** 1998. “DM-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer-Run Dependencies.” *The Journal of Finance*, 53: 219–265.

²⁰Favorable evidence for the intraday invariance hypothesis for a wide set of futures contracts and individual stocks is reported recently in Benzaquen, Donier and Bouchaud (2016).

- Andersen, Torben G., Tim Bollerslev, and Francis X. Diebold.** 2010. “Parametric and Nonparametric Measurements of Volatility.” In *Handbook of Financial Econometrics, Volume 1 — Tools and Techniques.*, ed. Yacine Ait-Sahalia and Lars Peter Hansen, 67–137. North Holland.
- Ané, Thierry, and Hélyette Geman.** 2000. “Order Flow, Transaction Clock, and Normality of Asset Returns.” *Journal of Finance*, 55(5): 2259–2284.
- Angel, James J.** 1997. “Tick Size, Share Prices, and Stock Splits.” *Journal of Finance*, 52(2): 655–681.
- Bae, Kyoungun, Albert S. Kyle, Eun Jung Lee, and Anna A. Obizhaeva.** 2014. “An Invariance Relationship in the Number of Buy-Sell Switching Points.” Working Paper, available at <http://dx.doi.org/10.2139/ssrn.2730770>.
- Benzaquen, Michael, Jonathan Donier, and Jean-Philippe Bouchaud.** 2016. “Unravelling the Trading Invariance Hypothesis.” Imperial College, London; Working Paper.
- Bertsimas, Dimitris, and Andrew W. Lo.** 1998. “Optimal Control of Liquidation Costs.” *Journal of Financial Markets*, 1: 1–50.
- Bollerslev, Tim, and Dan Jubinski.** 1999. “Equity Trading Volume and Volatility: Latent Information Arrivals and Common Long-Run Dependencies.” *Journal of Business & Economic Statistics*, 17: 9–21.
- Brennan, Michael, and Avandhar Subrahmanyam.** 1998. “The Determinants of Average Trade Size.” *The Journal of Business*, 71(1): 1–25.
- Brownlees, Christian T., Fabrizio Cipollini, and Giampiero M. Gallo.** 2011. “Intra-daily Volume Modeling and Prediction for Algorithmic Trading.” *Journal of Financial Econometrics*, 9: 489–518.
- Brownlees, Christian T., Fabrizio Cipollini, and Giampiero M. Gallo.** 2012. “Multiplicative Error Models.” In *Handbook of Financial Engineering and Econometrics: Volatility Models and Their Applications.*, ed. Luc Bauwens, Christian Hafner and Sebastian Laurent, 225–248. Wiley.
- Clark, Peter K.** 1973. “A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices.” *Econometrica*, 41(1): 135–155.
- Engle, Robert F.** 2002. “New Frontiers for ARCH Models.” *Journal of Applied Econometrics*, 17: 425–446.
- Gabaix, Xavier.** 2016. “Power Laws in Economics: An Introduction.” *Journal of Economic Perspectives*, 30(Winter): 185–206.

- Goldstein, Michael A., and Kenneth A. Kavajecz.** 2000. "Eighths, Sixteenths, and Market Depth: Changes in Tick Size and Liquidity Provision on the NYSE." *Journal of Financial Economics*, 56(1): 125–149.
- Harris, Lawrence.** 1987. "Transaction Data Tests of the Mixture of Distributions Hypothesis." *Journal of Financial and Quantitative Analysis*, 22(02): 127–141.
- Harris, Lawrence.** 1994. "Minimum Price Variations, Discrete Bid-ask Spreads, and Quotation Sizes." *Review of Financial Studies*, 7(1): 149–178.
- Jones, Charles M., Gautam Kaul, and Marc L. Lipson.** 1994. "Transactions, Volume and Volatility." *Review of Financial Studies*, 7(4): 631–651.
- Kirilenko, Andrei, Albert S. Kyle, Tugkan Tuzun, and Mehrdad Samadi.** 2014. "The Flash Crash: The Impact of High Frequency Trading on an Electronic Market." Working Paper, University of Maryland. Available at <http://dx.doi.org/10.2139/ssrn.1686004>.
- Kyle, Albert S., and Anna A. Obizhaeva.** 2016a. "Large Bets and Stock Market Crashes." Working Paper, available at <http://dx.doi.org/10.2139/ssrn.2023776>.
- Kyle, Albert S., and Anna A. Obizhaeva.** 2016b. "Market Microstructure Invariance: A Dynamic Equilibrium Model." Working Paper, University of Maryland.
- Kyle, Albert S., and Anna A. Obizhaeva.** 2016c. "Market Microstructure Invariance: Empirical Hypotheses." *Econometrica*, accepted for publication, available at <http://dx.doi.org/10.2139/ssrn.2722524>.
- Kyle, Albert S., Anna A. Obizhaeva, and Tugkan Tuzun.** 2016. "Microstructure Invariance in U.S. Stock Market Trades." Working Paper, available at <http://dx.doi.org/10.2139/ssrn.1107875>.
- Kyle, Albert S., Anna A. Obizhaeva, and Yajun Wang.** 2014. "Smooth Trading with Overconfidence and Market Power." Working Paper, available at <http://dx.doi.org/10.2139/ssrn.2423207>.
- Kyle, Albert S., Anna A. Obizhaeva, and Yajun Wang.** 2015. "Industrial Organization Implications of a One-Period Invariance Model." Working Paper.
- Kyle, Albert S., Anna A. Obizhaeva, Nitish R. Sinha, and Tugkan Tuzun.** 2014. "News Articles and the Invariance Hypothesis." Working Paper, available at <http://dx.doi.org/10.2139/ssrn.1786124>.
- Liesenfeld, Roman.** 2001. "A Generalized Bivariate Mixture Model for Stock Price Volatility and Trading Volume." *Journal of Econometrics*, 104: 141–178.
- Mandelbrot, Benoît, and Howard M. Taylor.** 1967. "On the Distribution of Stock Price Differences." *Operations Research*, 15(6): 1057–1062.

- Menkveld, Albert J., and Bart Z. Yueshen.** 2015. “The Flash Crash: a Cautionary Tale about Highly Fragmented Markets.” Working Paper, available at <http://dx.doi.org/10.2139/ssrn.2243520>.
- Moulton, Pamela C.** 2005. “You Can’t Always Get What You Want: Trade-size Clustering and Quantity Choice in Liquidity.” *Journal of Financial Economics*, 78(1): 89–119.
- Obizhaeva, Anna A., and Jiang Wang.** 2013. “Optimal Trading Strategy and Supply/Demand Dynamics.” *Journal of Financial Markets*, 16(1): 1–32.
- Schultz, Paul.** 2000. “Stock Splits, Tick Size, and Sponsorship.” *Journal of Finance*, 55(1): 429–450.
- Staffs of the CFTC and SEC.** 2010a. *Preliminary Findings Regarding the Market Events of May 6, 2010*. Report of the Staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues. May 18, 2010. “Preliminary Report”.
- Staffs of the CFTC and SEC.** 2010b. *Findings Regarding the Market Events of May 6, 2010*. Report of the Staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues. September 30, 2010. “Final Report”.
- Tauchen, George E., and Mark Pitts.** 1983. “The Price Variability-Volume Relationship on Speculative Markets.” *Econometrica*, 51: 485–505.

Appendix

A. High-Frequency Data Included in the Analysis

The set of high-frequency data used in our study consists of all one-minute observations available for the trading of the E-mini S&P 500 futures contract on the CME Group electronic Globex platform. As described in Section III, our sample covers January 4, 2008, to November 4, 2011. We use all observations within the continuous trading session, initiated each Sunday through Thursday, and ranging from 17:00 to 15:15 CT of the following day. Due to a few abbreviated trading sessions and loss of data dissemination, there are a few gaps in our series.

We drop the observations for the set of intervals listed below for which we have missing or incomplete records. First, we drop from the sample the entire trading day 03/12/2008 due to extremely low trading activity. Second, for the Asian regime, there is no trading on 01/02/2008, 12/26/2008, 01/02/2009, and 09/09/2009. These four gaps stretch into the European trading segments as well since observations are only available from 5:00 or 5:30 on these days. Third, we also have incomplete records for 01/22/2008 and 10/24/2008 during the European regime. Fourth, in the American regime, there are six days with reduced trading hours due to impending holidays: 07/03/2008, 11/28/2008, 12/24/2008, 11/27/2009, 12/24/2009, and 11/26/2010.

B. Inclusion of Price P

This section presents the findings from a set of representative regressions for time-series tests of intraday trading invariance when the price variation is included in the specification of the invariance hypothesis. It parallels the results reported in Table 5 of Section IV.C.

TABLE 7—TIME-SERIES REGRESSION OF $\log \frac{\sigma^2 P^2}{N}$ ONTO $\log Q$

Panel A: 3 Regimes						
	Nobs	c	β	se(c)	se(β)	\bar{R}^2
Asia	965	9.71	-0.85	0.062	0.041	0.310
Europe	969	9.21	-0.89	0.070	0.035	0.403
America	969	7.88	-0.61	0.115	0.045	0.157
Pooled	2903	11.09	-1.83	0.028	0.014	0.859

Panel B: 4 Years						
	Nobs	c	β	se(c)	se(β)	\bar{R}^2
2008	748	10.85	-1.66	0.054	0.027	0.833
2009	754	11.24	-2.02	0.044	0.022	0.917
2010	756	11.66	-2.11	0.051	0.024	0.909
2011	645	10.94	-1.71	0.053	0.024	0.884
All	2903	11.09	-1.83	0.028	0.014	0.859

This table reports on OLS regressions of the form, $s_{d,r} - n_{d,r} + 2p_{d,r} = c + \beta \cdot q_{d,r} + u_{d,r}$. For each day, there are three observations corresponding to the three regimes $r = 1, 2, 3$. In Panel A, the coefficients, standard errors, and \bar{R}^2 statistics are estimated separately for each regime. In Panel B, the regressions are estimated separately for each calendar year and then for the whole sample. The last year ends on November 4, 2011.