An Introduction to Market Microstructure Invariance

Albert S. Kyle
University of Maryland

Anna A. Obizhaeva
New Economic School

Imperial College
May 2016

Lecture 1

Lecture 1: Plan

- In the first lecture we would like to start with introducing market microstructure invariance as a set of empirical hypotheses.
- We will then describe persuasive empirical evidence on validity of invariance hypotheses based on a large data set of portfolio transitions.
- At the end, we will discuss the practical implications of market microstructure invariance and simple operational formulas for arrival of bets, the distribution of bet sizes, and transaction costs that use just a couple of calibrated constants.

Lecture 1: Literature

- Albert S. Kyle and Anna A. Obizhaeva, 2016, "Market Microstructure Invariance: Empirical Hypotheses," accepted for publication at Econometrica.
- Mark Kritzman, Albert S. Kyle, and Anna A. Obizhaeva, 2014, "A Practitioners Guide to Market Microstructure Invariance".

Overview

Our goal is to explain how **order size**, **order frequency**, **market efficiency and trading costs** vary across time and stocks.

- We propose **market microstructure invariance** that generates predictions concerning variations of these variables.
- We develop a meta-model suggesting that invariance is ultimately related to granularity of information flow.
- Invariance relationships are tested using a data set of portfolio transitions and find a strong support in the data.
- Invariance implies simple formulas for order size, order frequency, market efficiency, market impact, and bid-ask spread as functions of observable volume and volatility.

Preview of Results: Bet Sizes

Our estimates imply that bets $|\tilde{X}|/V$ are approximately distributed as **a log-normal** with the log-variance of 2.53 and the number of bets per day γ is defined as $(W = V \cdot P \cdot \sigma)$,

$$\ln \gamma = \ln 85 + \frac{2}{3} \ln \left[\frac{W}{(0.02)(40)(10^6)} \right].$$

$$\ln \left[\frac{|\tilde{X}|}{V} \right] \approx -5.71 - \frac{2}{3} \cdot \ln \left[\frac{W}{(0.02)(40)(10^6)} \right] + \sqrt{2.53} \cdot N(0,1)$$

For a benchmark stock, there are 85 bets with the median size of 0.33% of daily volume. Buys and sells are symmetric.

Preview of Results: Transaction Costs

Our estimates imply two simple formulas for expected trading costs for any order of X **shares** and for any security. The linear and square-root specifications are:

$$C(X) = \left(\frac{W}{(0.02)(40)(10^6)}\right)^{-1/3} \frac{\sigma}{0.02}\left(\frac{2.50}{10^4} \cdot \frac{X}{0.01\,V}\left[\frac{W}{(0.02)(40)(10^6)}\right]^{2/3} + \frac{8.21}{10^4}\right).$$

$$C(X) = \left(\frac{W}{(0.02)(40)(10^6)}\right)^{-1/3} \frac{\sigma}{0.02}\left(\frac{12.08}{10^4} \cdot \sqrt{\frac{X}{0.01\,V}\left[\frac{W}{(0.02)(40)(10^6)}\right]^{2/3}} + \frac{2.08}{10^4}\right).$$

Trading Games

We think of trading a stock as playing a trading game:

- **Long-term traders** buy and sell shares to implement "bets."
- **Intermediaries with short-term strategies**–market makers, high frequency traders, and other arbitragers–clear markets.

The intuition behind a trading game was first described by Jack Treynor (1971). In that game informed traders, noise traders and market makers traded with each other.

Since managers trade many different stocks, we can think of them as playing many different trading games simultaneously.

MAIN IDEA: Trading Games Across Stocks Are Played in "Business Time."

Stocks are different in terms of their trading activity: dollar trading volume, volatility etc. Trading games look different across stocks only at first sight!

Our intuition is that trading games are the same across stocks, except for the length of time over which these games are played or the speed with which they are played.

"Business time" passes faster for more actively traded stocks.

Games Across Stocks

Only the speed with which business time passes varies as trading activity varies:

- **For active stocks** (high trading volume and high volatility), trading games are played at a **fast pace**, i.e. the length of trading day is small and business time passes quickly.
- **For inactive stocks** (low trading volume and low volatility), trading games are played at a **slow pace**, i.e. the length of trading day is large and business time passes slowly.

Reduced Form Approach

As a rough approximation, we assume that bets arrive according to a Poisson process with **bet arrival rate** γ bets per day and **bet size** with a distribution $\tilde{Q}$ shares, $E(\tilde{Q}) = 0$.

Both $\tilde{Q}$ and γ vary across stocks.

Bet Volume and Bet Volatility

We define **bet volume** $\bar{V} := E|\tilde{Q}| \cdot \gamma = V/(\zeta/2)$.

We define **bet volatility** $\bar{\sigma} := \psi \cdot \sigma$.

ζ is "intermediation multiplier" and ψ is "volatility multiplier". We might assume ζ and ψ are constant, e.g., $\zeta = 2$ and $\psi = 1$.

Market Microstructure Invariance-1

Business time passes at a rate proportional to bet arrival rate γ, which measures market "velocity."

"Market Microstructure Invariance" is the hypothesis that *the dollar distribution of these gains or losses is the same across all markets when measured in units of business time*, i.e., the distribution of the random variable

$$\tilde{I} := P \cdot \tilde{Q} \cdot \left(\frac{\bar{\sigma}}{\gamma^{1/2}}\right)$$

is invariant across stocks or across time.

Market Microstructure Invariance-2

"Market Microstructure Invariance" is also the hypothesis that *the dollar cost of risk transfers is the same function of their size across all markets, when size of risk transfer is measured in units of business time*, i.e., trading costs of a risk transfer of size $\tilde{I}$,

$$C_B(\tilde{I})$$

is invariant across stocks or across time.

Trading Activity

Stocks differ in their **"trading activity"** W, or a measure of gross risk transfer, defined as dollar volume adjusted for volatility:

$$\bar{W} = \bar{\sigma} \cdot P \cdot \bar{V} = \bar{\sigma} \cdot P \cdot E|\tilde{Q}| \cdot \gamma.$$

Observable trading activity is a product of unobservable number of bets γ and bet size $\bar{\sigma} \cdot P \cdot E|\tilde{Q}|$.

Key Results

Since $\tilde{I} := P \cdot \tilde{Q} \cdot [\bar{\sigma}/\gamma^{1/2}]$ and $\bar{W} = \bar{\sigma} \cdot P \cdot E|\tilde{Q}| \cdot \gamma$, we have

$$\bar{W} = \gamma^{3/2} \cdot \{E|\tilde{I}|\}.$$

Therefore

$$\gamma = \bar{W}^{2/3} \cdot \{E|\tilde{I}|\}^{-2/3}.$$

$$\frac{\tilde{Q}}{\bar{V}} \sim \bar{W}^{-2/3} \cdot \{E|\tilde{I}|\}^{-1/3} \cdot \tilde{I}.$$

Frequency increases twice as fast as size, as trading speeds up.

Key Results

Define **asset-specific measure of liquidity** L by

$$L := \frac{\bar{W}^{1/3}}{\bar{\sigma}} \cdot E|\tilde{I}|^{1/3} = \left[\frac{P\bar{V}}{\bar{\sigma}^2}\right]^{1/3} \cdot E|\tilde{I}|^{-1/3}.$$

Define **invariant average price impact function** $f(I)$ by

$$f(\tilde{I}) := [C_B(\tilde{I})/\bar{C}_B]/[|\tilde{I}|/E|\tilde{I}|].$$

Then the percentage cost of executing a bet $C(\tilde{Q})$ is

$$C(\tilde{Q}) = \frac{C_B(\tilde{I})}{P|\tilde{Q}|} = \bar{\sigma}\bar{W}^{-1/3} \cdot \{E|\tilde{I}|\}^{1/3} \cdot f(\tilde{I}) = \frac{1}{L} \cdot f(\tilde{I}).$$

Costs Functions

- Linear model: $f(\tilde{I}) := \iota\bar{\kappa} + \iota^2\bar{\lambda} \cdot |\tilde{I}|$, where $\iota := (E|\tilde{I}|)^{-1/3}$.

$$C(\tilde{Q}) = \bar{\sigma}\left[\bar{\kappa} \cdot \bar{W}^{-1/3} + \bar{\lambda} \cdot \bar{W}^{1/3} \cdot \frac{|\tilde{Q}|}{\bar{V}}\right].$$

- Sqrt model: $f(\tilde{I}) := \iota\bar{\kappa} + \iota^{3/2}\bar{\lambda} \cdot |\tilde{I}|^{1/2}$.

$$C(\tilde{Q}) = \bar{\sigma}\left[\bar{\kappa} \cdot \bar{W}^{-1/3} + \bar{\lambda} \cdot \left|\frac{\tilde{Q}}{\bar{V}}\right|^{1/2}\right].$$

A Benchmark Stock

Benchmark Stock - daily volatility $\sigma = 200$ bps, price $P^* = \$40$, volume $V^* = 1$ million shares. Trades over a calendar day:

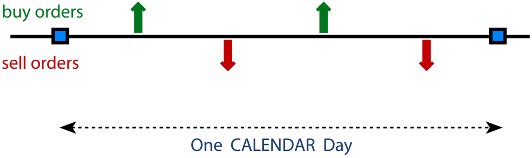

Arrival Rate $\gamma^* = 4$

Avg. Order Size $\bar{Q}^*$ **as fraction of** $V^* = 1/4$

Market Impact of 1/4 $V^* = 200$ bps / $4^{1/2} = 100$ bps

Market Microstructure Invariance - Intuition

Benchmark Stock with Volume V^*

$(\gamma^*, \quad \tilde{Q}^*)$

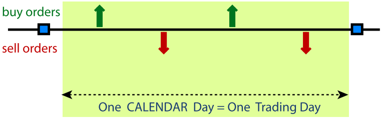

Avg. Order Size $\tilde{Q}^*$ as fraction of V^*

$= 1/4$

Market Impact of a Bet $(1/4 \ V^*)$

$= 200$ bps / $4^{1/2} = 100$ bps

Stock with Volume $V = \mathbf{8} \cdot V^*$

$(\gamma = \gamma^* \cdot 4, \quad \tilde{Q} = \tilde{Q}^* \cdot 2)$

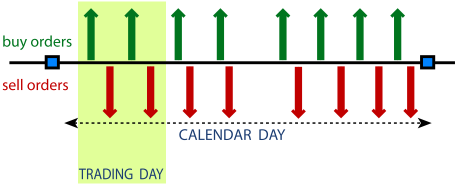

Avg. Order Size $\tilde{Q}$ as fraction of V

$= 1/16 = 1/4 \cdot \mathbf{8}^{-2/3}$

Market Impact of a Bet $(1/16 \ V)$

$= 200$ bps / $(4 \cdot \mathbf{8}^{2/3})^{1/2} = 50$ bps

$= 100$ bps $\cdot \mathbf{8}^{-1/3}$

Market Impact of $1/4 \ V$

$= 4 \cdot 50$ bps $= 100$ bps $\cdot \mathbf{8}^{1/3}$

Intuition: Invariance of Bets

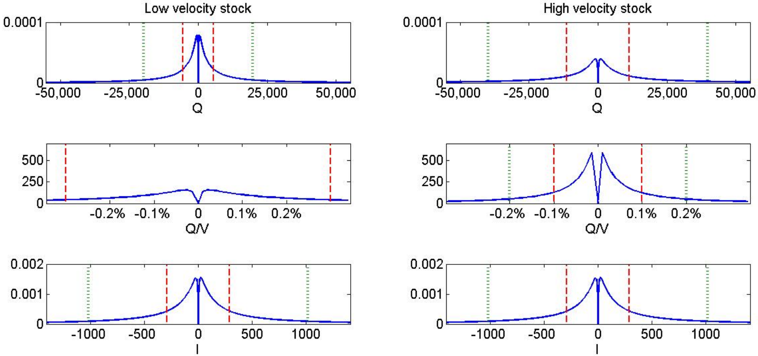

Distributions of $\tilde{Q}$ and $\tilde{Q}/V$ differ across stocks; distributions of $\tilde{I} := \tilde{Q} \cdot P \cdot \sigma \cdot \gamma^{-1/2}$ are the same.

Intuition: Invariance of Costs

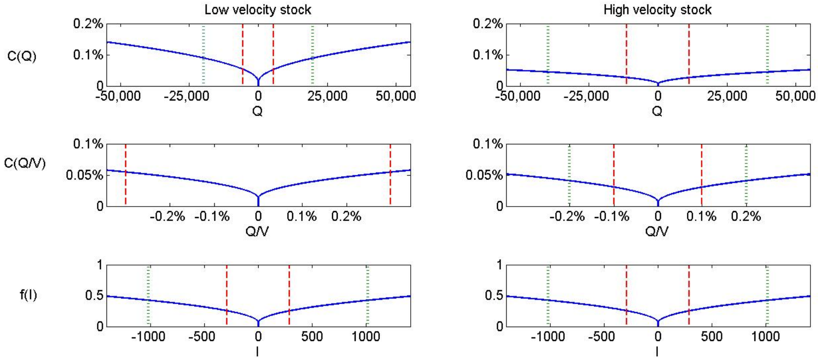

Percentage cost function $C(Q)$ and $C(Q/V)$ differ across stocks; function $f(I) := C(Q)/(1/L)$ for $I := Q \cdot P \cdot \sigma \cdot \gamma^{-1/2}$ are the same.

Example

Stock A: $\sigma = 0.02$, $P \cdot V = \$40 \cdot 10^6$.
100 bets per day (a bet per 4 minutes); median bet $\$100,000$.
Dollar risk transfer $P \cdot Q \cdot \sigma/\sqrt{\gamma} = \$200$ per unit of business time.

Stock B: $\sigma = 0.02$, $P \cdot V = 8 \cdot \$40 \cdot 10^6$.
$4 \cdot 100$ bets per day (a bet per 1 minutes); median bet $2 \cdot \$100,000$.
Dollar risk transfer $P \cdot Q \cdot \sigma/\sqrt{\gamma} = \$200$ per unit of business time.

The dollar cost of both bets is the same \$100. For the first bet the cost is 10 bps; for the second bet the cost is 5 bps.

Invariance Satisfies Irrelevance Principles

1. Modigliani-Miller Irrelevance: The trading game involving a financial security issued by a firm is **independent** of its **capital structure:**

- Stock Split Irrelevance,
- Leverage Irrelevance.

2. Time-Clock Irrelevance: The trading game is **independent** of the **time clock**. Transaction costs functions and illiquidity measure $1/L$ remain the same regardless of whether a researcher measures $\gamma, \bar{V}, \bar{\sigma}$, and $\bar{W}$ using different time horizons.

Invariance and Previous Literature

Microstructure invariance does not undermine or contradict other theoretical models of market microstructure. It builds a bridge from theoretical models to empirical tests of those models.

- **Theoretical models** usually suggest that order flow imbalances move prices, but do not provide a unified framework for mapping the theoretical concept of an order flow imbalance into empirically observed variables.
- **Empirical tests** often use "wrong" proxies for unobserved order imbalances such as volume or square root of volume.

Microstructure invariance is **a modeling principle** making it possible to test theoretical models empirically.

Example: Invariance and Kyle (1985)

Kyle (1985) and other models imply a linear price impact formula

$$\lambda = \frac{\sigma_V}{\sigma_U}$$

where σ_V is the standard deviation of dollar price change per share resulting from price impact, and σ_U is the standard deviation of "order imbalances".

$$\lambda := \frac{\bar{\sigma}}{\gamma^{1/2} \cdot (E\tilde{Q}^2)^{1/2}} = \frac{\bar{\sigma}}{\bar{V}} \cdot \bar{W}^{1/3} \cdot [E\{|\tilde{I}|^2\}]^{-1/2}[E\{|\tilde{I}|\}]^{-2/3}.$$

In data, calibrate the constant $[E\{|\tilde{I}|^2\}]^{-1/2}[E\{|\tilde{I}|\}]^{-2/3}$.

Testing - Portfolio Transition Data

The empirical implications of the three proposed models are tested using a proprietary dataset of **portfolio transitions**.

- Portfolio transition occurs when an old (legacy) portfolio is replaced with a new (target) portfolio during replacement of fund management or changes in asset allocation.
- Our data includes 2,550+ portfolio transitions executed by a large vendor of portfolio transition services over the period from 2001 to 2005.
- Dataset reports executions of 400,000+ orders with average size of about 4% of ADV.

Portfolio Transitions and Trades

We use the data on **transition orders** to examine which model makes the most reasonable assumptions about how the **size of trades** varies with **trading activity**.

Super-Cloud: Log of Order Size vs. Log of Trading Activity

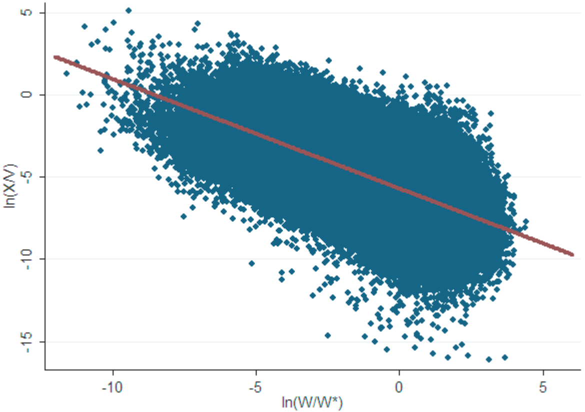

The figure shows $\ln\left[\frac{X_i}{V_i}\right]$ as function of $\ln\left[\frac{W_i}{W^*}\right]$. All observations of line up along the line with the slope of $-2/3$.

Distribution of Order Sizes

Microstructure invariance predicts that **distributions** of order sizes X, adjusted for differences in trading activity W, are the same across different stocks:

$$\ln\left(\frac{|\tilde{Q}|}{V} \cdot \left[\frac{W}{W^*}\right]^{2/3}\right).$$

We compare distributions across **10 volume/5 volatility groups**.

Distributions of Order Sizes

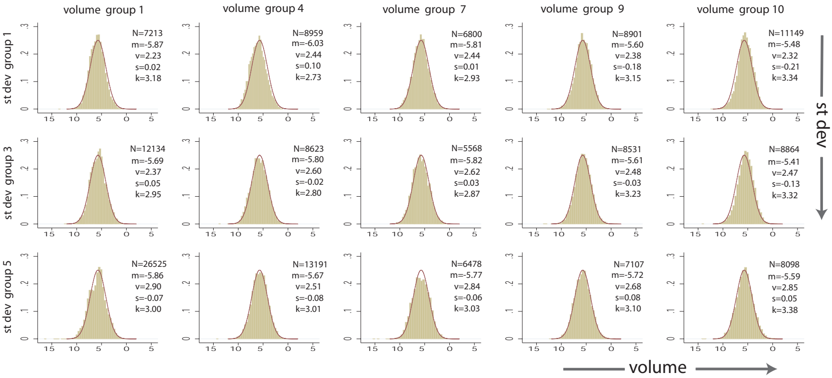

Microstructure invariance works well for **entire distributions** of order sizes. These distributions are approximately **log-normal** with log-variance of 2.53.

Log-Normality of Order Size Distributions

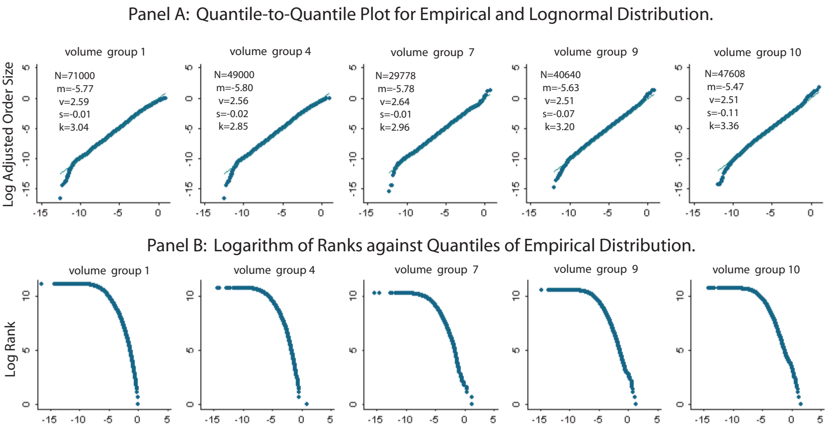

Microstructure invariance works well for **entire distributions** of order sizes. These distributions are approximately **log-normal**.

Tests for Orders Size - Design

In regression equation that relates **trading activity** W and the trade size $\tilde{Q}$, proxied by **a transition order of** X **shares**, as a fraction of average daily volume V:

$$\ln\left[\frac{X_i}{V_i}\right] = ln[\bar{q}] + a_0 \cdot \ln\left[\frac{W_i}{W_*}\right] + \tilde{\epsilon}$$

Microstructure Invariance predicts $a_0 = -2/3$.

The variables are scaled so that $\bar{q}$ is (assuming log-normal distribution) the median size of liquidity trade as a fraction of daily volume for **a benchmark stock** with daily standard deviation of 2%, price of \$40 per share, trading volume of 1 million shares per day, ($W_* = 0.02 \cdot 40 \cdot 10^6$).

Tests for Order Size: Results

	All	NYSE Buy	NYSE Sell	NASDAQ Buy	NASDAQ Sell
$\ln\left[\bar{q}\right]$	**-5.67**	-5.68	-5.63	-5.75	-5.65
	(0.017)	(0.023)	(0.018)	(0.035)	(0.032)
α_0	**-0.62**	**-0.63**	**-0.59**	**-0.71**	**-0.59**
	(0.009)	(0.011)	(0.008)	(0.019)	(0.015)

- **Microstructure Invariance:** $a_0 = -2/3$.

Why Coefficients for Sells Different from Buys

- Since asset managers are "long only," buys are related to current value of W, while sells are related to value of W when stocks were bought.
- Since increases in W result from positive returns, higher values of W are correlated with higher past returns.
- Implies sell coefficients smaller in absolute value than buy coefficients, consistent with empirical results.
- Adding lagged returns or lagged trading activity W may improve results.

Percentiles Tests for Order Size: Results

	p1	p5	p25	p50	p75	p95	p99
$\ln\left[\bar{q}\right]$	-9.37	-8.31	-6.73	-5.66	-4.59	-3.05	-2.05
	(0.008)	(0.006)	(0.004)	(0.003)	(0.004)	(0.006)	(0.009)
α_0	-0.65	-0.64	-0.61	-0.62	-0.61	-0.64	-0.63
	(0.005)	(0.003)	(0.002)	(0.002)	(0.002)	(0.003)	(0.005)

- **Microstructure Invariance:** $a_0 = -2/3$.

Tests for Orders Size - R^2

	All	NYSE Buy	NYSE Sell	NASDAQ Buy	NASDAQ Sell
Unrestricted Specification: $\alpha_0 = -2/3$					
R^2	0.3229	0.2668	0.2739	0.4318	0.3616
Restricted Specification: $b_1 = b_2 = b_3 = b_4 = 0$					
R^2	0.3167	0.2587	0.2646	0.4298	0.3542
Microstructure Invariance: $\alpha_0 = -2/3$, $b_1 = b_2 = b_3 = b_4 = 0$					
R^2	0.3149	0.2578	0.2599	0.4278	0.3479

$$\ln\left[\frac{X_i}{V_i}\right] = \ln\left[\bar{q}\right] - \alpha_0 \cdot \ln\left[\frac{W_i}{W*}\right] + b_1 \cdot \ln\left[\frac{\sigma_i}{0.02}\right] + b_2 \cdot \ln\left[\frac{P_{0,i}}{40}\right] + b_3 \cdot \ln\left[\frac{V_i}{10^6}\right] + b_4 \cdot \ln\left[\frac{\nu_i}{1/12}\right] + \tilde{\epsilon}.$$

Tests for Orders Size - Summary

Microstructure Invariance predicts: An increase of one percent in trading activity W leads to a decrease of $2/3$ of one percent in bet size as a fraction of daily volume (for constant returns volatility).

Results: The estimates provide strong support for microstructure invariance. The coefficient predicted to be -2/3 is estimated to be -0.62.

Discussion:

- The assumptions made in our model match the data economically.
- F-test rejects our model statistically because of small standard errors.
- Invariance explains data for buys better than data for sells.
- Estimating coefficients on P, V, σ, ν improves R^2 very little compared with imposing coefficient value of $-2/3$.

Portfolio Transitions and Trading Costs

We use data on the **implementation shortfall** of portfolio transition trades to test predictions of the three proposed models concerning how **transaction costs,** both market impact and bid-ask spread, vary with **trading activity.**

Portfolio Transitions and Trading Costs

"Implementation shortfall" is the difference between actual trading prices (average execution prices) and hypothetical prices resulting from "paper trading" (price at previous close).

There are **several problems** usually associated with using implementation shortfall to estimate transactions costs. Portfolio transition orders avoid most of these problems.

Problem I with Implementation Shortfall

Implementation shortfall is a **biased estimate** of transaction costs when it is based on price changes and executed quantities, because these quantities themselves are often correlated with price changes in a manner which biases transactions costs estimates.

Example A: Orders are often canceled when price runs away. Since these non-executed, high-cost orders are left out of the sample, we would underestimate transaction costs.

Example B: When a trader places an order to buy stock, he has in mind placing another order to buy more stock a short time later.

For **portfolio transitions**, this problem does not occur: Orders are not canceled. The timing of transitions is somewhat exogenous.

Problems II with Implementation Shortfall

The second problem is **statistical power**.

Example: Suppose that 1% ADV has a transactions cost of 20 bps, but the stock has a volatility of 200 bps. Order adds only 1% to the variance of returns. A properly specified regression will have an R squared of 1% only!

For **portfolio transitions**, this problem does not occur: Large and numerous orders improve statistical precision.

Tests For Transaction Costs - Design

In the regression specification that relates **trading activity** W and **implementation shortfall C** for a transition order for X shares:

$$\mathbb{I}_{BS,i} \cdot C(X_i) \cdot \frac{(0.02)}{\sigma_i} = a \cdot R_{mkt} \cdot \frac{(0.02)}{\sigma_i} + \mathbb{I}_{BS,i} \cdot \left[\frac{W_i}{W^*}\right]^{\alpha} \cdot C^*(I_i) + \epsilon_i.$$

Microstructure invariance predicts that $\alpha = -1/3$ and function $C^*(I)$ does not vary across stocks and time. Function $C^*(I) = L^* \cdot f(I)$ quantifies the trading costs for a benchmark stock.

- Implementation shortfall is adjusted for market changes.
- Implementation shortfall is adjusted for differences in volatility.

Percentiles Tests for Quoted Spread: Results

		NYSE		NASDAQ	
	All	Buy	Sell	Buy	Sell
$\ln\left[k^*/(40 \cdot 0.02)\right]$	**-3.07**	-3.09	-3.08	-3.04	-3.04
	(0.008)	(0.008)	(0.008)	(0.013)	(0.012)
α_1	**-0.35**	**-0.31**	**-0.32**	**-0.40**	**-0.39**
	(0.003)	(0.003)	(0.003)	(0.004)	(0.004)

- **Microstructure Invariance:** $a_1 = -1/3$.

$$\ln\left[\frac{\kappa_i}{P_{0,i}\sigma_i}\right] = \ln\left[\frac{k^*}{40 \cdot 0.02}\right] + \alpha_1 \cdot \ln\left[\frac{W_i}{W^*}\right] + \tilde{\epsilon}.$$

Results Related to Quoted Spread

Regression of log of spread on log of trading activity W:

- Predicted coefficient is $-1/3$.
- Estimated coefficient is -0.35, being different for NYSE (-0.31)and for NASDAQ (-0.40).

Using quoted spread rather than implicit realized spread cost in transactions cost regression, we get estimated coefficient of 0.71, with puzzling variation across buys (0.61) and sells (0.75).

Tests For Market Impact and Spread: Results

	NYSE		NASDAQ		
	All	Buy	Sell	Buy	Sell
a	0.66	0.63	0.62	0.76	0.78
	(0.013)	(0.016)	(0.016)	(0.037)	(0.036)
$1/2\bar{\lambda}^* \times 10^4$	**10.69**	12.08	9.56	12.33	9.34
	(1.376)	(2.693)	(2.254)	(2.356)	(2.686)
z	**0.57**	0.54	0.56	0.44	0.63
	(0.039)	(0.056)	(0.062)	(0.051)	(0.086)
α_2	**-0.32**	**-0.40**	**-0.33**	**-0.41**	**-0.29**
	(0.015)	(0.037)	(0.029)	(0.035)	(0.037)
$1/2\bar{\kappa}^* \times 10^4$	**1.77**	-0.27	1.14	0.77	3.55
	(0.837)	(2.422)	(1.245)	(4.442)	(1.415)
α_3	**-0.49**	**-0.37**	**-0.50**	**0.53**	**-0.44**
	(0.050)	(1.471)	(0.114)	(1.926)	(0.045)

- **Microstructure Invariance:** $\alpha_2 = -1/3, \alpha_3 = -1/3$.

$$\mathbb{I}_{BS,i} \cdot C(X_i) \cdot \frac{(0.02)}{\sigma_i} = a \cdot R_{mkt} \cdot \frac{(0.02)}{\sigma_i} + \frac{\bar{\lambda}^*}{2}\mathbb{I}_{BS,i} \cdot \left[\frac{\phi I_i}{0.01}\right]^z \cdot \left[\frac{W_i}{W^*}\right]^{\alpha_2} + \frac{\bar{\kappa}^*}{2}\mathbb{I}_{BS,i} \cdot \left[\frac{W_i}{W^*}\right]^{\alpha_3} + \tilde{\epsilon}.$$

Discussion

- Estimated coefficient $a = 0.66$ suggests that most orders are executed within one day.
- In a non-linear specification, α_3 is often different from predicted -1/3, but spread cost $\bar{\kappa}$ is insignificant.
- Scaled cost functions are non-linear with the estimated exponent $z = 0.57$.
- Buys have higher price impact $\bar{\lambda}^*$ than sells, since buys may be more informative whereas price reversals after sells makes their execution cheaper.

Tests for Transaction Costs - R^2

	All	NYSE Buy	NYSE Sell	NASDAQ Buy	NASDAQ Sell
	Unrestricted Specification, 12 Degrees of Freedom: $\alpha_2 = \alpha_3 = -1/3$				
R^2	0.1016	0.1121	0.1032	0.0957	0.0944
	Restricted Specification: $\beta_1 = \beta_2 = \beta_3 = \beta_4 = \beta_5 = \beta_6 = \beta_7 = \beta_8 = 0$				
R^2	0.1010	0.1118	0.1029	0.0945	0.0919
	Microstructure Invariance, SQRT Model: $z = 1/2, \beta_1 = \beta_2 = \beta_3 = \beta_4 = \beta_5 = \beta_6 = \beta_7 = \beta_8 = 0, \alpha_2 = \alpha_3 = -1/3$				
R^2	0.1007	0.1116	0.1027	0.0941	0.0911
	Microstructure Invariance, Linear Model: $z = 1, \beta_1 = \beta_2 = \beta_3 = \beta_4 = \beta_5 = \beta_6 = \beta_7 = \beta_8 = 0, \alpha_2 = \alpha_3 = -1/3$				
R^2	0.0991	0.1102	0.1012	0.0926	0.0897

$$\mathbb{I}_{BS,i} \cdot C(X_i) \cdot \frac{(0.02)}{\sigma_i} = a \cdot R_{mkt} \cdot \frac{(0.02)}{\sigma_i} + \frac{\bar{\lambda}^*}{2} \mathbb{I}_{BS,i} \cdot \left[\frac{\phi I_i}{0.01}\right]^z \cdot \left[\frac{W_i}{W*}\right]^{\alpha_2} \cdot \frac{\sigma_i^{\beta_1} \cdot P_{0,i}^{\beta_2} \cdot V_i^{\beta_3} \cdot \nu_i^{\beta_4}}{(0.02)(40)(10^6)(1/12)}$$

$$+ \frac{\bar{\kappa}^*}{2} \mathbb{I}_{BS,i} \cdot \left[\frac{W_i}{W*}\right]^{\alpha_3} \cdot \frac{\sigma_i^{\beta_5} \cdot P_{0,i}^{\beta_6} \cdot V_i^{\beta_7} \cdot \nu_i^{\beta_8}}{(0.02)(40)(10^6)(1/12)} + \tilde{\epsilon}.$$

Tests for Trading Costs - Summary

Microstructure Invariance predicts: An increase of one percent in trading activity W leads to a decrease of $1/3$ of one percent in transaction costs (for constant returns volatility).

Results: The estimates provide strong support for microstructure invariance. The coefficient predicted to be -1/3 is estimated to be -0.32.

Discussion:

- Invariance matches the data economically.
- F-test rejects invariance statistically because of small standard errors.
- Price impact cost is better described by a non-linear function with exponent of 0.57.
- Estimating coefficients on P, V, σ, ν improves R^2 very little comparing with imposing coefficient of $-1/3$, especially comparing to a square root model.

Transactions Costs Across Volume Groups

For each of **10 volume groups/100 order size groups**, we estimate dummy coefficients from regression:

$$\mathbb{I}_{BS,i} \cdot C(X_i) \cdot \frac{(0.02)}{\sigma_i} = a \cdot R_{mkt} \cdot \frac{(0.02)}{\sigma_i} + \mathbb{I}_{BS,i} \cdot \left[\frac{W_i}{W^*}\right]^{-1/3} \cdot \sum_{j=1}^{100} \mathbb{I}_{i,j,k} \cdot c^*_{k,j}.$$

- Indicator variable $\mathbb{I}_{i,j,k}$ is one if ith order is in the kth volume groups and jth size group.
- The dummy variables $c^*_{k,j}, j = 1, ..100$ track the shape of scaled transaction costs function $C^*(I)$ for kth volume group.

If invariance holds, then all estimated functions should be the same across volume groups.

Transactions Costs Across Volume Groups

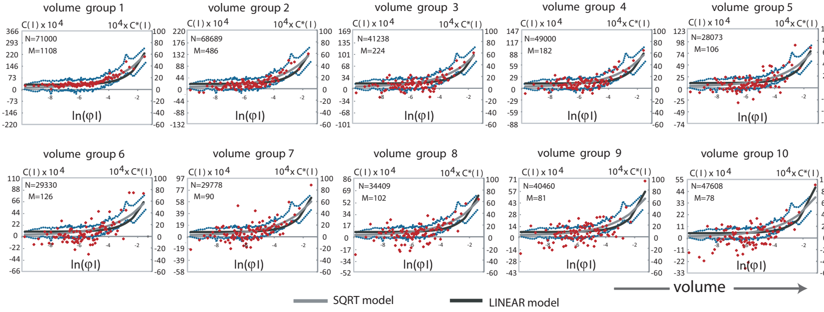

For each of 10 volume groups, 100 estimated dummy variables $c^*_{k,j}, j = 1, ..100$ track scaled cost functions $C^*(I)$ for a benchmark stock on the left axis. Actual costs functions $C(I)$ are on the right axis. Group 1 contains stocks with the lowest volume. Group 10 contains stocks with the highest volume. The volume thresholds are 30th, 50th, 60th, 70th, 75th, 80th, 85th, 90th, and 95th percentiles for NYSE stocks.

Transactions Costs by Percentiles of $\tilde{I}$

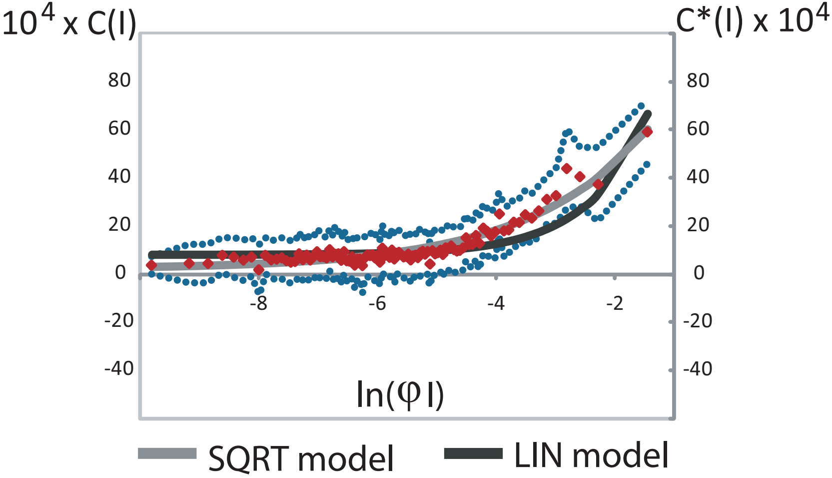

Invariance of Cost Functions - Discussion

- Cost functions scaled by $\sigma W^{-1/3}$ with argument X scaled by $W^{2/3}/V$ seem to be stable across volume groups.

- The estimates are more "noisy" in higher volume groups, since transitions are usually implemented over one calendar day, i.e., over longer horizons in business time for larger stocks.

- The square-root specification fits the data slightly better than the linear specification, particularly for large orders in size bins from 90th to 99th.

- The linear specification fits better costs for very large orders in active stocks.

Calibration: Bet Sizes

Our estimates imply that portfolio transition orders $|\tilde{X}|/V$ are approximately distributed as **a log-normal** with the log-variance of 2.53 and the number of bets per day γ is defined as,

$$\ln \gamma = \ln 85 + \frac{2}{3} \ln \left[\frac{W}{(0.02)(40)(10^6)} \right].$$

$$\ln \left[\frac{|\tilde{X}|}{V} \right] \approx -5.71 - \frac{2}{3} \cdot \ln \left[\frac{W}{(0.02)(40)(10^6)} \right] + \sqrt{2.53} \cdot N(0,1)$$

For a benchmark stock, there are 85 bets with the median size of 0.33% of daily volume. Buys and sells are symmetric.

Calibration: Transactions Cost Formula

Our estimates imply two simple formulas for expected trading costs for any order of X **shares** and for any security. The linear and square-root specifications are:

$$C(X) = \left(\frac{W}{(0.02)(40)(10^6)}\right)^{-1/3} \frac{\sigma}{0.02}\left(\frac{2.50}{10^4}\cdot\frac{X}{0.01\,V}\left[\frac{W}{(0.02)(40)(10^6)}\right]^{2/3}+\frac{8.21}{10^4}\right).$$

$$C(X) = \left(\frac{W}{(0.02)(40)(10^6)}\right)^{-1/3} \frac{\sigma}{0.02}\left(\frac{12.08}{10^4}\cdot\sqrt{\frac{X}{0.01\,V}\left[\frac{W}{(0.02)(40)(10^6)}\right]^{2/3}}+\frac{2.08}{10^4}\right).$$

More Practical Implications

- **Trading Rate:** If it is reasonable to restrict trading of the benchmark stock to say 1% of average daily volume, then a smaller percentage would be appropriate for more liquid stocks and a larger percentage would be appropriate for less liquid stocks.
- **Components of Trading Costs:** For orders of a given percentage of average daily volume, say 1%, bid-ask spread is a relatively larger component of transactions costs for less active stocks, and market impact is a relatively larger component of costs for more active stocks.
- **Comparison of Execution Quality:** When comparing execution quality across brokers specializing in stocks of different levels of trading activity, performance metrics should take account of nonlinearities documented in our paper.

Calibration: Bet Size and Trading Activity

For **a benchmark stock** with $40 million daily volume and 2% daily returns standard deviation, empirical results imply:

- Median bet size is $132,500 or 0.33% of daily volume.
- Average bet size is $469,500 or 1.17% of daily volume.
- Benchmark stock has about 85 bets per day.
- Order imbalances are 38% of daily volume.
- Half price impact is 2.50 and half spread is 8.21 basis points.
- Expected cost of a bet is about $2,000.

Invariance allows to extrapolate these estimates to other assets.

Calibration: Implications of Log-Normality for Volume and Volatility

Standard deviation of log of bet size is $2.53^{1/2}$ implies:

- a one-standard-deviation increase in bet size is a factor of about 4.90.
- 50% of trading volume generated by largest 5.39% of bets.
- 50% of returns variance generated by largest 0.07% of bets (linear model).

“Time Change” Literature

“Time change” is the idea that a larger than usual number of independent price fluctuations results from business time passing faster than calendar time.

- Mandelbrot and Taylor (1967): Stable distributions with kurtosis greater than normal distribution implies infinite variance for price changes.
- Clark (1973): Price changes result from log-normal with time-varying variance, implying finite variance to price changes.
- Econophysics: Gabaix et al. (2006); Farmer, Bouchard, Lillo (2009). Right tail of distribution might look like a power law.
- **Microstructure invariance**: Kurtosis in returns results from rare, very large bets, due to high variance of log-normal. Caveat: Large bets may be executed very slowly, e.g., over weeks.

Conclusions

- Predictions of microstructure invariance largely hold in portfolio transitions data for equities.
- We conjecture that invariance predictions can be found to hold as well in other datasets and may generalize to other markets and other countries.
- We conjecture that market frictions such as wide tick size and minimum round lot sizes may result in deviations from the invariance predictions. Invariance provides a benchmark for measuring the importance of those frictions.
- Microstructure invariance has numerous implications.

Lecture 2

Lecture 2: Plan

- In the second lecture we will talk about how to derive invariance relationships by combining dimensional analysis with Modigliani-Miller invariance.
- We will also talk about applications for various markets as well as market events—such as the U.S. stock market crashes in 1929 and 1987, Flash crash in May 2010, the flash rally in U.S. Treasuries market in October 2014, the crash in the Russian currency market in December 2014, and the Chinese stock market crash in 2015—and discuss why market microstructure invariance can help to explain large price changes during these historical episodes.

Lecture 2: Literature

- Albert S. Kyle and Anna A. Obizhaeva, 2016, “Market Microstructure Invariance and Dimensional Analysis”
- Albert S. Kyle and Anna A. Obizhaeva, 2016, “Large Bets and Stock Market Crashes”

Implication for Market Crashes

Order of 5% of daily volume is “normal” for a typical stock. Order of 5% of daily volume is “unusually large” for the market.

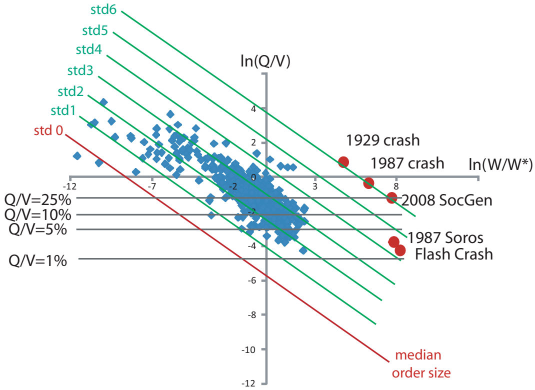

Conventional intuition that order equal to 5% of average daily volume will not trigger big price changes in indices is wrong!

Calibration of Market Crashes

	Actual	Predicted Invariance	Predicted Conventional	%ADV	%GDP
1929 Market Crash	**25%**	**44.35%**	1.36%	241.52%	1.136%
1987 Market Crash	**32%**	**16.77%**	0.63%	66.84%	0.280%
1987 Soros's Trades	**22%**	**6.27%**	0.01%	2.29%	0.007%
2008 SocGén Trades	**9.44%**	**10.79%**	0.43%	27.70%	0.401%
2010 Flash Crash	**5.12%**	**0.61%**	0.03%	1.49%	0.030%

Table shows the actual price changes, predicted price changes, orders as percent of average daily volume and GDP, and implied frequency.

Discussion

- **Price impact predicted by invariance is large and similar to actual price changes.**
- **The financial system in 1929 was remarkably resilient.** The 1987 portfolio insurance trades were equal to about 0.28% of GDP and triggered price impact of 32% in cash market and 40% in futures market. The 1929 margin-related sales during the last week of October were equal to 1% of GDP. They triggered price impact of 24% only.

Discussion - Cont'd

- **Speed of liquidation magnifies short-term price effects.** The 1987 Soros trades and the 2010 flash-crash trades were executed rapidly. Their actual price impact was greater than predicted by microstructure invariance, but followed by rapid mean reversion in prices.

- **Market crashes happen too often.** The three large crash events were approximately 6 standard deviation bet events, while the two flash crashes were approximately 4.5 standard deviation bet events. Right tail appears to be fatter than predicted. The true standard deviation of underlying normal variable is not 2.53 but 15% bigger, or far right tail may be better described by a power law.

Early Warning System

Early warning systems may be useful and practical. Invariance can be used as a practical tool to help quantify the systemic risks which result from sudden liquidations of speculative positions.

Dimensional Analysis

Invariance hypotheses can be derived based on the "dimensional analysis," in a manner similar to Kolmogorov's laws in theory of turbulence:

Notation:

- $1/L =$ "illiquidity" (unitless measure of transactions costs)
- $P =$ Price (dollars per share, e.g., \$40)
- $V =$ Volume (shares per day, e.g., one million shares per day)
- $\sigma^2 =$ Returns Variance (unitless per day, e.g., 2% per day squared $= 0.04$)
- $C = =$ Cost of a bet (dollars, e.g., \$2,000 per bet)

Suppose $1/L$ is log-linear in other variables:

$$\frac{1}{L} = P^{\delta_1} \cdot V^{\delta_2} \cdot \sigma^{\delta_3} \cdot C^{\delta_4}. \tag{1}$$

Dimensional Analysis (continued)

Cancellation of units requires $\delta := \delta_4 = \delta_3/2 = -\delta_1 = -\delta_2$:

$$\frac{1}{L} = \left(\frac{C \cdot \sigma^2}{P \cdot V} \right)^{\delta}. \tag{2}$$

What is δ? If stock levered up by a factor of two, Modigliani-Miller equivalence suggests that P halves, σ doubles, and $1/L$ doubles. This implies $\delta = 1/3$. We obtain

$$\frac{1}{L} = \left(\frac{C \cdot \sigma^2}{P \cdot V} \right)^{1/3}. \tag{3}$$

Compare with Amihud's measure, which is similar to

$$\frac{1}{L_{Amihud}} = \frac{C \cdot \sigma}{P \cdot V}, \tag{4}$$

where time units do not cancel and MM equivalence does not hold!

Invariance-Implied Liquidity Measures

$\gamma =$ "Velocity": Total dollars in expected transactions costs per day:

$$\gamma \cdot E\{C(\tilde{I})\} \propto \gamma, \qquad \gamma = W^{2/3} = [P \cdot V \cdot \sigma]^{2/3}$$

$L_\sigma =$ "Risk Liquidity": Cost of transferring a risk:

$$E\left\{\frac{C(\tilde{I})}{|P\tilde{Q}\sigma|}\right\} \propto \frac{1}{L_\sigma} \cdot f(\tilde{I}), \qquad L_\sigma := \mathrm{const} \cdot W^{1/3} = \mathrm{const} \cdot [P \cdot V \cdot \sigma]^{1/3}$$

$L_\$ =$ "Dollar Liquidity": Cost of Converting Asset to Cash (basis points):

$$E\left\{\frac{C(\tilde{I})}{|P\tilde{Q}|}\right\} \propto \frac{1}{L_\$} \cdot f(\tilde{I}), \qquad L_\$:= \mathrm{const} \cdot \frac{W^{1/3}}{\sigma} = \mathrm{const} \cdot \left[\frac{P \cdot V}{\sigma^2}\right]^{1/3}$$

Linear and Square Root Costs

Linear Costs: Suppose $f(\tilde{I}) = f(W^{2/3}\tilde{Q}/V)$ linear in $\tilde{I}$.
$\chi =$ "Market Temperature" defined by Derman (2002).:

$$E\left\{\frac{C(\tilde{I})}{|P\tilde{Q}|}\right\} = \frac{\sigma}{W^{1/3}} \cdot f(W^{2/3}\frac{\tilde{Q}}{V}) \propto \chi \cdot \frac{|\tilde{Q}|}{V},$$

$$\chi := \sigma \cdot \gamma^{1/2} = \sigma \cdot W^{1/3} = \sigma^{4/3} \cdot (P \cdot V)^{1/3}$$

Square Root Costs: Suppose $f(\tilde{I}) = f(W^{2/3}\tilde{Q}/V)$ square root in $\tilde{I}$

$$E\left\{\frac{C(\tilde{I})}{|P\tilde{Q}|}\right\} = \sigma \cdot W^{-1/3} \cdot f(W^{2/3}\tilde{Q}/V) \propto \sigma \cdot \frac{\tilde{Q}}{V}$$

Only σ and bet size as fraction of volume $|\tilde{Q}|/V$ matters!

Lecture 3

Lecture 3: Plan

- In the final lecture we will discuss why scaling laws can be derived in the context of more conventional microstructure equilibrium models.
- We will also discuss other empirical findings supporting predictions of market microstructure invariance. The examples include the evidence from our studies of the U.S. stock market, the E-mini S&P 500 futures market, the Korean stock market, the Russian stock market, and Thomson-Reuters news articles data. Invariance explains a substantial part of cross-sectional and time-series variations and provides a useful benchmark for studying various market frictions.

Lecture 3: Literature

- Albert S. Kyle and Anna A. Obizhaeva, 2016, "Market Microstructure Invariance: A Dynamic Equilibrium Model"
- Albert S. Kyle, Anna A. Obizhaeva, and Tugkan Tuzun, 2016, "Microstructure Invariance in the U.S. stock market trades"
- Albert S. Kyle, Anna A. Obizhaeva, Nitish Sinha, and Tugkan Tuzun, 2011, "News Articles and the Invariance Hypothesis"
- Torben Andersen, Oleg Bondarenko, Albert S. Kyle, Anna A. Obizhaeva, and Tugkan Tuzun, 2016, "Intraday Trading Invariance in the E-mini S&P 500 Futures Market"
- Kyoung-hun Bae, Alber S. Kyle, Eun Jung Lee, and Anna A. Obizhaeva, 2014, "An Invariance Relationship in the Number of Buy-Sell Switching Points"

A Structural Model

We outline a dynamic infinite-horizon model of trading, from which various invariance relationships are derived results.

- **Informed traders** face given costs of acquiring information of given precision, then place informed bets which incorporate a given fraction of the information into prices.
- **Noise traders** place bets which turn over a constant fraction of the stocks float, mimicking the size distribution of bets placed by informed trades.
- **Market makers** offer a residual demand curve of constant slope, lose money from being "run over" by informed bets, but make up the losses from trading costs imposed on informed and noise traders.

Fundamental Value

- The unobserved "fundamental value" of the asset follows an exponential martingale:

$$F(t) := \exp[\sigma_F \cdot B(t) - \tfrac{1}{2} \cdot \sigma_F^2 \cdot t],$$

where $B(t)$ follows standardized Brownian motion with $\mathrm{var}\{B(t + \Delta t) - B(t)\} = \Delta t$. $F(t)$ follows a martingale.

Market Prices

- The price changes as informed traders and noise traders arrive in the market and anonymously place bets.
- Risk neutral market makers set the market price $P(t)$ as the conditional expectation of the fundamental value $F(t)$ given a history of the “bet flow”.
- $\bar{B}(t)$ is the market's conditional expectation of $B(t)$ based on observing the history of prices; the error $B(t) - \bar{B}(t)$ has a normal distribution with variance denoted $\Sigma(t)/\sigma_F^2$.
- The price is the best estimate of fundamental value; the price has a martingale property:

$$P(t) = \exp[\sigma_F \cdot \bar{B}(t) + \tfrac{1}{2} \cdot \Sigma(t) - \tfrac{1}{2} \cdot \sigma_F^2 \cdot t].$$

Pricing Accuracy

- Pricing accuracy is defined as $\Sigma(t) = var\{\log[F(t)/P(t)]\}$; market is more efficient when $\Sigma^{1/2}$ is smaller.
- $\Sigma^{-1/2}$ is Fischer Black's measure of market efficiency: He conjectures *"almost all markets are efficient"* *in the sense that "price is within a factor 2 of value" at least* 90% *of the time.*

Pricing Accuracy - Intuition

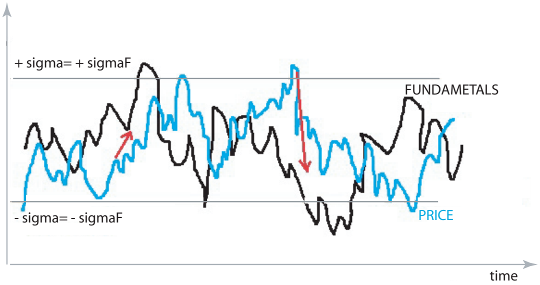

- Pricing accuracy is defined as $\Sigma(t) = var\{\log[F(t)/P(t)]\}$; the market is more efficient when $\Sigma^{1/2}$ is smaller.
- Fama says a market is "efficient" if all information is appropriately reflected in price (prices follow a martingale), even if very little information is available and prices are not very accurate, i.e., $\Sigma^{1/2}$ is large.

Pricing Accuracy

- $\Sigma^{-1/2}$ is Fischer Black's measure of market efficiency: He conjectures *"almost all markets are efficient" in the sense that "price is within a factor 2 of value" at least* 90% *of the time.* In mathematical terms, $\Sigma^{1/2} = \ln(2)/1.64 = 0.42$.

- In time units, Σ/σ^2 is the number of years by which the informational content of prices lags behind fundamental value, e.g., if $\sigma = 0.35$ and $\Sigma^{1/2} = \ln(2)/1.64$, then prices are about $(\ln(2)/1.64)^2/0.35^2 \approx 1.50$ years "behind" fundamental value.

Informed Traders

- Informed traders arrive randomly in the market at rate $\gamma_I(t)$.
- Each informed trader observes one private signal $\tilde{i}(t)$ and places one and only one bet, which is executed by trading with market makers.

$$\tilde{i}(t) := \tau^{1/2} \cdot \Sigma(t)^{-1/2} \cdot \sigma_F \cdot [B(t) - \bar{B}(t)] + \tilde{Z}_I(t),$$

where τ measures the precision of the signal and $\tilde{Z}_I(t) \sim N(0, 1)$. $\mathrm{var}\{\tilde{i}(t)\} = 1 + \tau \approx 1$.

Informed Traders

- An informed trader updates his estimate of $B(t)$ from $\bar{B}(t)$ to $\bar{B}(t) + \Delta \bar{B}_I(t)$. Assuming τ is small,

$$\Delta \bar{B}_I(t) \approx \tau^{1/2} \cdot \Sigma(t)^{1/2} / \sigma_F \cdot \tilde{i}(t).$$

- If the signal value were to be fully incorporated into prices, then the dollar price change would be equal to

$$E\{F(t) - P(t) \,|\, \Delta \bar{B}_I(t)\} \approx P(t) \cdot \sigma_F \cdot \Delta \bar{B}_I(t).$$

- Only a fraction θ of the "fully revealing" impact is incorporated into prices ($\lambda(t)$ is price impact), i.e.,

$$\tilde{Q}(t) = \theta \cdot \lambda(t)^{-1} \cdot P(t) \cdot \sigma_F \cdot \Delta \bar{B}_I(t).$$

Profits of Informed Traders

- An informed trader's expected "paper trading" profits are

$$\bar{\pi}_I(t) := E\{[F(t) - P(t)] \cdot \tilde{Q}(t)\} = \frac{\theta \cdot P(t)^2 \cdot \sigma_F^2 \cdot E\{\Delta \bar{B}_I(t)^2\}}{\lambda(t)}.$$

- His expected profits net of costs conditional on $\Delta \bar{B}_I(t)$ are

$$E\{[F(t) - P(t)] \cdot \tilde{Q}(t) - \lambda(t) \tilde{Q}(t)^2\} = \frac{\theta(1-\theta) P(t)^2 \sigma_F^2 \cdot \Delta \bar{B}_I(t)^2}{\lambda(t)}.$$

- $\theta = 1/2$ maximizes the expected profits of the risk-neutral informed trader. We assume $0 < \theta < 1$ to accommodate possibility of informed traders being risk averse and information can be leaked.

Noise Traders

- Noise traders arrive at an endogenously determined rate $\gamma_U(t)$.
- Each noise trader places one bet which mimics the size distribution of an informed bet, even though it contains no information, i.e., $\tilde{I}(t) = \tilde{Z}_U(t) \sim N(0, 1 + \tau) \approx N(0, 1)$.
- Noise traders turn over a constant fraction η of shares outstanding N. The expected share volume $V(t)$ and total number of bets per day $\gamma(t) := \gamma_I(t) + \gamma_U(t)$ satisfy

$$\gamma_U(t) \cdot E\{|\tilde{Q}(t)|\} = \eta \cdot N, \qquad \gamma(t) \cdot E\{|\tilde{Q}(t)|\} = V(t).$$

Transaction Costs

- Both informed traders and noise traders incur transactions costs. The unconditional expected costs are

$$\bar{C}_B(t) := \lambda(t) \cdot E\{\tilde{Q}(t)^2\} = \frac{\theta^2 \cdot P(t)^2 \cdot \sigma_F^2 \cdot E\{\Delta \bar{B}(t)^2\}}{\lambda(t)}.$$

- Illiquidity $1/L(t)$ is defined as the expected cost of executing a bet in basis points:

$$1/L(t) := \bar{C}_B(t)/E\{|P(t) \cdot \tilde{Q}(t)|\}.$$

Break-Even Conditions - Intuition

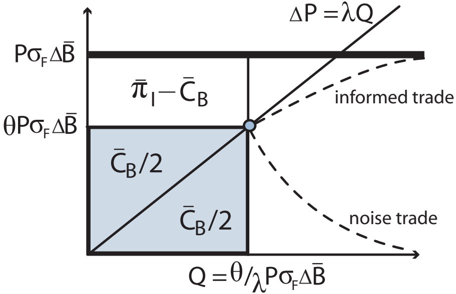

There is price continuation after an informed trade and mean reversion after a noise trade. The losses on trading with informed traders are equal to total gains on trading with noise traders, $\gamma_I \cdot (\bar{\pi}_I - \bar{C}_B) = \gamma_U \cdot \bar{C}_B$.

Break-Even Condition For Market Maker

- The equilibrium level of costs must allow market makers to break even.
- The expected dollar price impact costs that market makers expect to collect from all traders must be equal to the expected dollar paper trading profits of informed traders:

$$(\gamma_I(t) + \gamma_U(t)) \cdot \bar{C}_B(t) = \gamma_I \cdot \bar{\pi}_I(t).$$

Break-Even Condition for Informed Traders

- The break-even condition for informed traders yields the rate at which informed traders place bets $\gamma_I(t)$.
- The expected paper trading profits from trading on a signal $\bar{\pi}_I(t)$ must equal the sum of expected transaction costs $\bar{C}_B(t)$ and the exogenously constant cost of acquiring private information denoted c_I:

$$\bar{\pi}_I(t) = \bar{C}_B(t) + c_I.$$

Market Makers and Market Efficiency

- Zero-profit, risk neutral, competitive market makers set prices such that the price impact of anonymous trades reveals on average the information in the order flow. The average impact of a bet must satisfy

$$\lambda(t) \cdot \tilde{Q}(t) = \frac{\gamma_I(t)}{\gamma_I(t) + \gamma_U(t)} \cdot \lambda(t) \cdot \tilde{Q}(t) \cdot \frac{1}{\theta} + \frac{\gamma_U(t)}{\gamma_I(t) + \gamma_U(t)} \cdot 0.$$

- The ratio of informed traders to noise traders then turns out to be equal to the exogenous constant θ. The turnover rate is constant.

$$\frac{\gamma_I(t)}{\gamma_I(t) + \gamma_U(t)} = \theta, \quad V = \eta \cdot N/(1 - \theta).$$

Diffusion Approximation

- As a result of each bet, market makers update their estimate of $\bar{B}(t)$ by $\Delta \bar{B}(t)$.
- A trade is informed with probability θ and, if informed, incorporates a fraction θ of its information content into prices, leading to an adjustment in $\bar{B}(t)$ of

$$\Delta \bar{B}(t) = \theta \tau^{1/2} \Sigma(t)^{1/2} \sigma_F^{-1} \cdot \left(\tau^{1/2} \Sigma(t)^{-1/2} \sigma_F [B(t) - \bar{B}(t)] + \tilde{Z}_I(t) \right)$$

- A trade is uninformed with probability $1 - \theta$ and, if uninformed, adds noise to $\bar{B}(t)$ of

$$\Delta \bar{B}(t) = \theta \tau^{1/2} \Sigma(t)^{1/2} \sigma_F^{-1} \cdot \tilde{Z}_U(t).$$

Diffusion Approximation

- When the arrival rate of bets $\gamma(t)$ per day is sufficiently large, the diffusion approximation for the dynamics of the estimate $\bar{B}(t)$ can be written as

$$d\bar{B}(t) = \gamma(t)\cdot\theta^2\cdot\tau\cdot[B(t)-\bar{B}(t)]\cdot dt+\gamma(t)^{1/2}\cdot\theta\cdot\tau^{1/2}\cdot\Sigma(t)^{1/2}\cdot\sigma_F^{-1}\cdot dZ(t).$$

The first term corresponds to the information contained in informed signals which arrive at rate $\theta \cdot \gamma(t)$. The second term corresponds to the noise contained in all bets arriving at rate $\gamma(t)$.

Equilibrium Price Process

- Define

$$\sigma(t) := \theta \cdot \tau^{1/2} \cdot \Sigma(t)^{1/2} \cdot \gamma(t)^{1/2}.$$

- By applying Ito's lemma,

$$\frac{dP(t)}{P(t)} = \frac{1}{2} \cdot [\Sigma'(t) - \sigma_F^2 + \sigma(t)^2] \cdot dt + \sigma_F \cdot d\bar{B}(t).$$

- Market efficiency implies that $P(t)$ must follow a martingale:

$$\frac{d\Sigma(t)}{dt} = \sigma_F^2 - \sigma(t)^2.$$

Price Process

- Since in the equilibrium,

$$\frac{dP(t)}{P(t)} = \sigma(t) \cdot d\bar{Z}(t).$$

The process $\bar{Z}(t)$ is a standardized Brownian motion under the market's filtration and $\sigma(t)$ is the measure of returns volatility.

Resiliency

- The difference $B(t) - \bar{B}(t)$ follows the mean-reverting process,

$$d[B(t) - \bar{B}(t)] = -\frac{\sigma(t)^2}{\Sigma(t)} \cdot [B(t) - \bar{B}(t)] \cdot dt + dB(t) - \frac{\sigma(t)}{\sigma_F} \cdot dZ(t).$$

- Market resiliency $\rho(t)$ be the mean reversion rate at which pricing errors disappear

$$\rho(t) = \frac{\sigma(t)^2}{\Sigma(t)}.$$

Holding returns volatility constant, resiliency is greater in markets with higher pricing accuracy.

Invariance Theorem - 1

Assume the cost c_I *of generating a signal is an invariant constant and let* $m := E\{|\tilde{i}(t)|\}$ *define an additional invariant constant.*

Then, the invariance conjectures hold: The dollar risk transferred by a bet per unit of business time is a random variable with an invariant distribution $\tilde{I}$*, and the expected cost of executing a bet* $\bar{C}_B$ *is constant:*

$$\tilde{I}(t) := P(t) \cdot \tilde{Q}(t) \cdot \frac{\sigma(t)}{\gamma(t)^{1/2}} = \bar{C}_B \cdot \tilde{i}(t).$$

$$\bar{C}_B = c_I \cdot \theta/(1-\theta).$$

$$C_B(t) = \frac{1}{\bar{c}_B} \cdot I^2, \qquad where \qquad I(t) \equiv \frac{P(t) \cdot Q \cdot \sigma(t)}{\gamma(t)^{1/2}}.$$

Invariance Theorem -2

The number of bets per day $\gamma(t)$, their size $\tilde{Q}(t)$, liquidity $L(t)$, pricing accuracy $\Sigma(t)^{-1/2}$, and market resiliency $\rho(t)$ are related to price $P(t)$, share volume $V(t)$, volatility $\sigma(t)$, and trading activity $W(t) = P(t) \cdot V(t) \cdot \sigma(t)$ by the following invariance relationships:

$$\gamma(t) = \left(\frac{\lambda(t) \cdot V(t)}{\sigma(t)P(t)m}\right)^2 = \left(\frac{E\{|\tilde{Q}(t)|\}}{V(t)}\right)^{-1} = \frac{(\sigma(t)L(t))^2}{m^2} = \frac{\sigma(t)^2}{\theta^2\tau\Sigma(t)} = \frac{\rho(t)}{\theta^2\tau} = \left(\frac{W(t)}{m\bar{C}_B}\right)^{2/3}.$$

Arrival Rate—Impact—Bet Size—Liquidity—"Efficiency"—Resilience—Activity

τ is the precision of a signal, θ is the fraction of information $\tilde{i}(t)$ incorporated by an informed trade. The price follows a martingale with stochastic returns volatility $\sigma(t) := \theta \cdot \tau^{1/2} \cdot \Sigma(t)^{1/2} \cdot \gamma(t)^{1/2}$.

Proof

The proof is based on the solution of the system of four equations:

- Volume condition: $\gamma(t) \cdot E\{|\tilde{Q}(t)|\} = V(t)$
- Market resiliency $\bar{c}_B = \lambda(t) \cdot E\{\tilde{Q}^2(t)\}$,
- Volatility condition: $\gamma(t) \cdot \lambda(t)^2 \cdot E\{\tilde{Q}(t)^2\} = P(t)^2 \cdot \sigma(t)^2$,
- Moments ratio: $m = \frac{E\{|\tilde{Q}(t)|\}}{[E\{\tilde{Q}(t)^2\}]^{1/2}}$.

One can think of $\gamma(t)$, $\lambda(t)$, $E\{\tilde{Q}(t)^2\}$, and $E\{|\tilde{Q}(t)|\}$ as unknown variables to be solved for in terms of known variables $V(t)$, $\bar{c}_B$, $P(t)$, and $\sigma(t)$.

Discussion

- Trading activity $W(t)$ and its components—prices $P(t)$, share volume V, and returns volatility $\sigma(t)$—are a "macroscopic" quantities, which are easy to estimate.
- The bet arrival rate $\gamma(t)$, bet size $\tilde{Q}(t)$, the average cost of a bet $1/L(t)$, pricing accuracy $\Sigma(t)^{1/2}$, and resiliency $\rho(t)$ are "microscopic" quantities, which are difficult to estimate.

Invariance relationships allow to infer microscopic quantities from macroscopic quantities ($\bar{C}_B$, m, and $\tau \cdot \theta^2$ are just constants).

Discussion

- The assumption that distinct bets result from distinct pieces of private information implies a particular level of granularity for both signals and bets.
- The invariance of bet sizes and their cost rely on the assumption that cost of a private signal c_I and the shape of the distribution of signals m are constant (c_I can be replaced by productivity-adjusted wage of a finance professional).
- The invariance of pricing accuracy and resiliency requires stronger assumptions: the informativeness of a bet $\tau \cdot \theta^2$ is constant.
- The model is motivated by the time series properties of a single stock as its market capitalization changes, but it can apply cross-sectionally across different securities.

Robustness of Assumptions

- Our structural model makes numerous restrictive assumptions. The empirical results we are about to describe are not consistent with the "linear-normal" assumptions of the model.
- The size of unsigned bets closely fits a log-normal distribution, not a normal distribution. A linear price impact model predicts market impact costs reasonably well, but a square root model of price predicts impact costs better.
- We conjecture that it should be possible to modify our structural model to accommodate those issues.
- The model is to be interpreted as a "proof of concept" consistent with the interpretation that the invariance hypotheses might apply more generally.

Robustness of Assumptions

- Our structural model makes numerous restrictive assumptions. The empirical results we are about to describe are not consistent with the “linear-normal” assumptions of the model.
- The size of unsigned bets closely fits a log-normal distribution, not a normal distribution. A linear price impact model predicts market impact costs reasonably well, but a square root model of price predicts impact costs better.
- We conjecture that it should be possible to modify our structural model to accommodate those issues.
- The model is to be interpreted as a “proof of concept” consistent with the interpretation that the invariance hypotheses might apply more generally.

Smooth Trading Model

Albert S. Kyle, Anna A. Obizhaeva, and Yajun Wang. "Smooth Trading with Overconfidence and Market Power." Available at SSRN 2423207 (2013).

Spirit of the meta-model can be implemented as a model similar to smooth trading model.

Leads to invariance relationships.

Summary of Invariance Papers

Market microstructure invariance is formulated for bets:

- Albert S. Kyle and Anna A. Obizhaeva, 2016, "Market Microstructure Invariance: Empirical Hypotheses"
- Albert S. Kyle and Anna A. Obizhaeva, 2016, "Market Microstructure Invariance: A Dynamic Equilibrium Model"

Its extension "intraday trading invariance" is formulated for trades:

- Albert S. Kyle, Anna A. Obizhaeva, and Tugkan Tuzun, 2016, "Microstructure Invariance in the U.S. stock market trades"
- Torben Andersen, Oleg Bondarenko, Albert S. Kyle, Anna A. Obizhaeva, and Tugkan Tuzun, 2016, "Intraday Trading Invariance in the E-mini S&P 500 Futures Market"
- Kyoung-hun Bae, Alber S. Kyle, Eun Jung Lee, and Anna A. Obizhaeva, 2014, "An Invariance Relationship in the Number of Buy-Sell Switching Points"

Other stuff:

- Albert S. Kyle, Anna A. Obizhaeva, Nitish Sinha, and Tugkan Tuzun, 2011, "News Articles and the Invariance Hypothesis"
- Albert S. Kyle and Anna A. Obizhaeva, 2016, "Dimensional Analysis and Market Microstructure Invariance"

OLS Estimates of Number of Trades in TAQ

Trading game invariance seems to **work** in TAQ **before 2001,** subject to market frictions (Kyle, Obizhaeva and Tuzun (2010)).

	1993/2008	**1993/2000**	**2001/2008**
α	7.10	6.15	8.04
	(0.076)	(0.021)	(0.060)
a_γ	**0.74**	**0.69**	**0.79**
	(0.004)	(0.001)	(0.004)
Adj-R^2	0.92	0.91	0.94
# Obs	5,801	6,698	4,914

Evidence From TAQ Dataset Before 2001

Here are trade-weighted distributions, NYSE stocks, 1993.

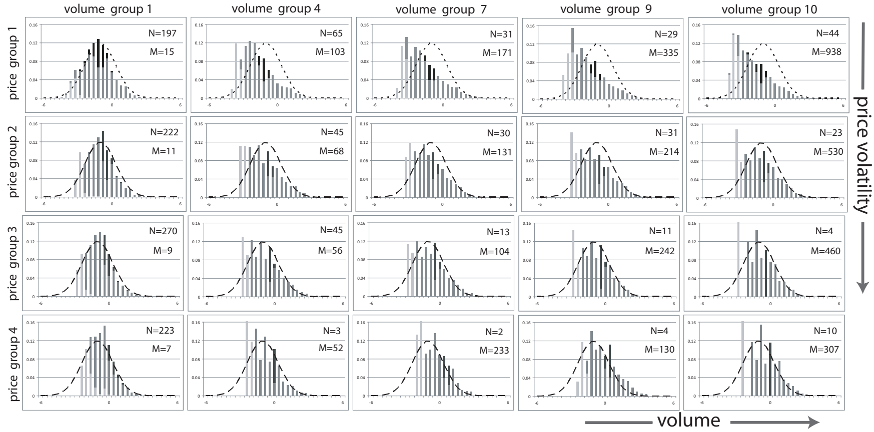

Evidence From TAQ Dataset Before 2001

Here are volume-weighted distributions, NYSE stocks, 1993.

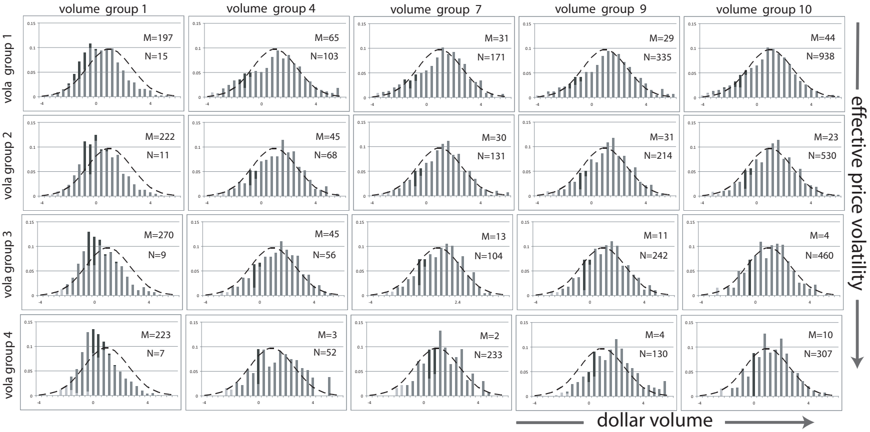

Evidence From TAQ Dataset After 2001

Trading game invariance is **hard to test** in TAQ **after 2001.**

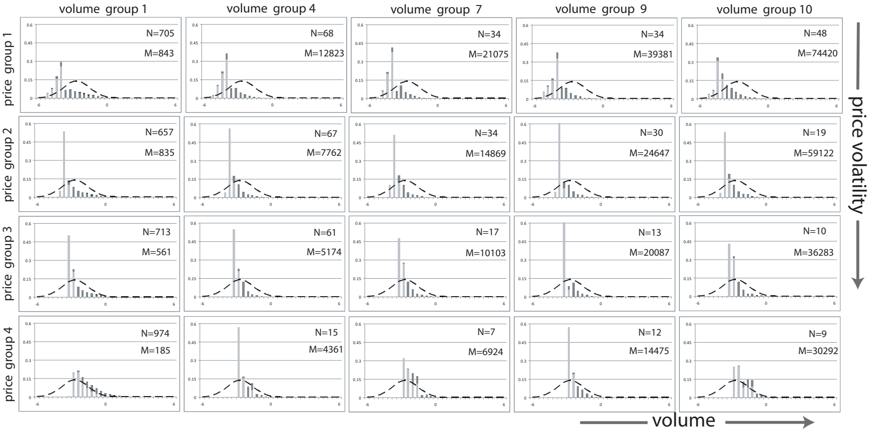

Intraday Patterns for S&P500 E-mini Futures

Intraday patterns for volume, volatility, number of trades, and average trade size (Andersen, Bondarenko, Kyle, Obizhaeva (2014))

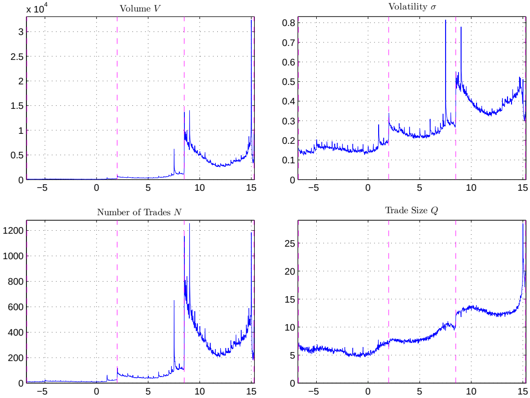

Intraday Patterns for S&P500 E-mini Futures

Log of number of trades on log trading activity, by day and one-minute time interval (2008-2011). Predicted coeff. is 2/3. The fitted line is $\ln(N_{dt}) = -3.7415 + \mathbf{0.661} \cdot \ln(V_{dt} \cdot P_{dt} \cdot \sigma_{dt})$.

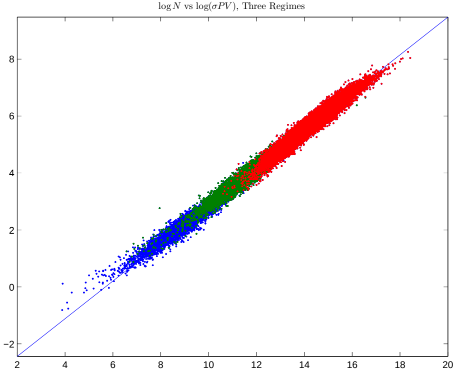

Switching Points in South Korean Market

Log of buy-sell switching points on log trading activity, by stock and month. Predicted coeff. is 2/3. The fitted line is $\ln(S_{it}) = 11.156 + \mathbf{0.675} \cdot \ln(W_{it}/W^*)$. (Bae, Kyle, Lee, Obizhaeva (2014))

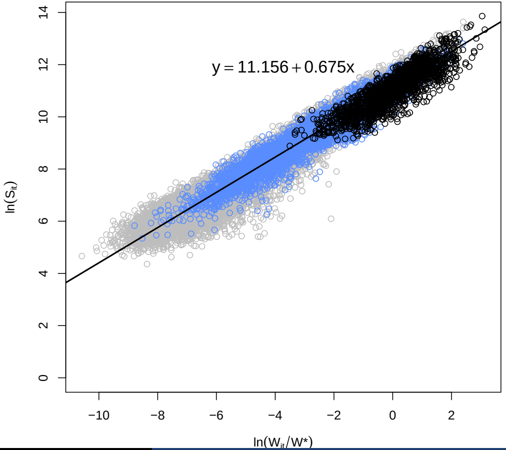

Time Series of Monthly Coefficients

The estimates of slope from monthly regressions of the log of aggregate number of switching points on the log of trading activity fluctuate over time between 0.64 and 0.72.

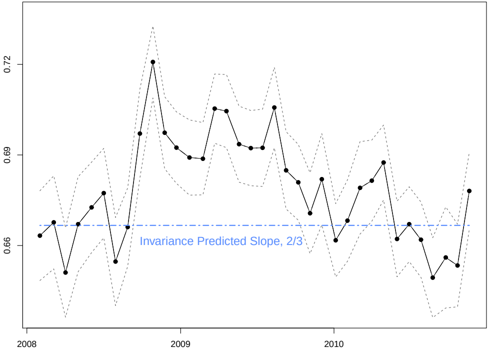

Different Types of Traders

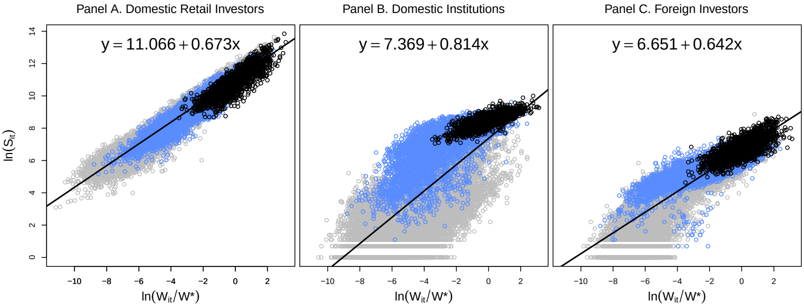

Trading by retail investors most closely satisfy predictions.

Number of News Articles: Distribution.

Figure shows the distribution of the number of news articles per month for ten volume groups across twelve news bins with 0, 1, 2, 3-4, 5-8, 9-16, 17-32, 33-64, 65-128, 129-256, 257-512, 513-1024 news items per month.

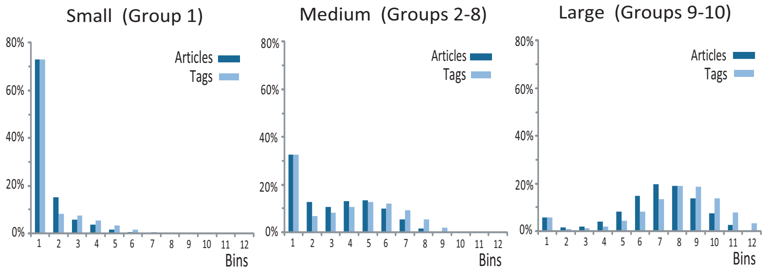

We examine whether the market microstructure invariance explains the differences in the distributions of arrivals of news articles.

Invariance Predication

As the time clock speeds up, the rate of information flow increases proportionately. When trading activity W increases by one percent, the rate of information flow speeds up by a two-third of one percent,

$$\mu(W) \sim W^{2/3}.$$

Data

The news data, collected in NewsScope dataset, is provided by the Thomson Reuters firm.

- Sample: NewsScope dataset provided by the Thomson Reuters firm to its client.
- Period: January 2003 - December 2008.
- Coverage: most U.S. stocks.
- There are 1.4 million news articles and 3.4 million news tags mentioning a particular topic code. The average number of firms in a given month is 3,820.

Poisson Model

The distribution of the number of news items $N_{t,i}$ about stock i in month t has the density function,

$$f(N_{t,i}|W_{t,i}) = \frac{e^{-\mu(W_{t,i})} \times \mu(W_{t,i})^{N_{t,i}}}{N_{t,i}!},$$

where the arrival rate of news items $\mu_{t,i}$ is a non-linear function,

$$\mu(W_{t,i}) = e^{\eta+\alpha \cdot \ln\left[\frac{\bar{W}_{t,j}}{W^*}\right]}.$$

Stocks have the same expected number of news items $\mu(W)$, equal to its variance, controlling for W. This assumption may be too restrictive.

Negative Binomial Model

NB model allows the Poisson arrival rate to vary randomly, even for firms with the same level of trading activity. This variation is modeled with a continuous mixture of the Poisson distributions and the gamma distribution,

$$\mu(W_{t,i}) = e^{\eta + \alpha \cdot \ln\left[\frac{\bar{W}_{t,j}}{W^*}\right]} \cdot \tilde{G}_{t,i}(\beta).$$

The Gamma variable $\tilde{G}_{t,i}(\beta)$ has the mean equal to one and its variance to be equal to β.

Poisson model if $\beta = 0$. Data is over-dispersed if $\beta > 0$.

The Poisson Regression Estimates

	News Articles		News Tags	
	All	Thomson-Reuters	All	Thomson-Reuters
η	2.11	2.19	2.78	2.85
	(0.044)	(0.037)	(0.049)	(0.044)
α	**0.81**	**0.78**	**0.86**	**0.84**
	(0.007)	(0.007)	(0.008)	(0.010)
$\log(L)$	-16,590	-15,722	-33,249	-31,570

- **Model of Microstructure Invariance:** $\alpha = 2/3$.
- Model of Invariant Bet Frequency: $\alpha = 0$.
- Model of Invariant Bet Size: $\alpha = -1$.

The Negative Binomial Regression Estimates

	News Articles		News Tags	
	All	Thomson-Reuters	All	Thomson-Reuters
η	2.01	2.08	2.66	2.73
	(0.036)	(0.028)	(0.037)	(0.030)
α	**0.68**	**0.65**	**0.71**	**0.66**
	(0.024)	(0.018)	(0.025)	(0.019)
β	2.05	1.63	3.17	2.49
	(0.218)	(0.120)	(0.325)	(0.170)
$\log(L)$	-17,170	-6,900	-8,584	-8,289

- **Model of Microstructure Invariance:** $\alpha = 2/3$.
- Model of Invariant Bet Frequency: $\alpha = 0$.
- Model of Invariant Bet Size: $\alpha = -1$.

News Articles and Invariance

Data on the number of Reuters news items N is consistent with trading game invariance (Kyle, Obizhaeva, Ranjan, and Tuzun (2010)).

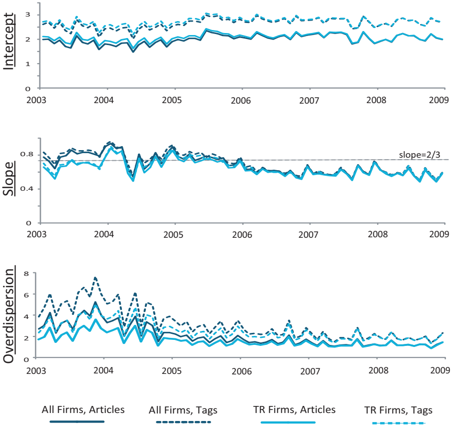

Calibration

The news articles distribution for a stock with share price P, expected daily volume V shares, and expected daily volatility σ is well described by the Negative Binomial model with the expected arrival rate:

$$\mu(W) = e^{1.97} \cdot \left[\frac{V \cdot P \cdot \sigma}{0.02 \cdot 40 \cdot 10^6}\right]^{2/3} \cdot \tilde{G}(2.11).$$

- For a benchmark stock, there are about 7 news articles per month.
- For a small stock, there are about 0.87 news articles per month.
- For a large stock, there are about 40 news articles per month.