Centre for
Economic
and Financial
Research
at
New Economic
School

January 2017

Dimensional Analysis and Market Microstructure

Albert S. Kyle
Anna A. Obizhaeva

Working Paper No 234

CEFIR / NES Working Paper series

Dimensional Analysis and Market Microstructure Invariance.

Albert S. Kyle[1], Anna A. Obizhaeva[2]

[1] *University of Maryland, College Park, MD 20742, USA*
[2] *New Economic School, Moscow, Skolkovo, 143026, Russia*

Market microstructure is the subfield of finance and econophysics[1] which studies how prices result from the process of trading securities. Large trades move prices[2] and incur trading costs. Here we combine dimensional analysis, leverage neutrality, and a principle of market microstructure invariance to derive scaling laws which express transaction costs functions, bid-ask spreads, bet sizes, number of bets, and other financial variables in terms of trading volume and volatility. For example, market liquidity is proportional to the cube root of the ratio of dollar volume to return variance. We illustrate the scaling by showing that bid-ask spreads in Russian stocks indeed scale with the cube root. In addition to being of interest to risk managers and traders, these scaling laws provide scientific benchmarks for evaluating controversial issues related to high frequency trading, market crashes, and liquidity measurement as well as guidelines for designing policies in the aftermath of financial crisis.

Physics researchers obtain powerful results by using dimensional analysis[3] to reduce the dimensionality of problems. For example, Kolmogorov proposed a simple dimensional analysis argument to derive his "5/3-law" for the energy distribution in a turbulent fluid.[4] His law describes the relationship between the energy spectrum, energy flow, and wave lengths. Similar analysis can be used to infer the size and number of molecules in a mole of gas or the size of the explosive energy in an atomic blast from measurable large-scale physical quantities. While finance and economics are fields which respect consistency of units—e.g., by maintaining a distinction between stocks and flows—dimensional analysis is not generally used to derive non-obvious but powerful results.

In financial markets, institutional investors trade by implementing speculative "bets" which move prices. A bet is a decision to buy or sell a quantity of institutional size. Traders execute bets by dividing them into separate orders, shredding the orders into small pieces, and executing numerous smaller quantities over time. For securities, the time frame of execution may be minutes, hours, days, or weeks. Here we use dimensional analysis to derive scaling laws for transaction costs functions, the width of bid-ask spreads, the size distribution of bets, the speed of bet execution, the natural minimum price fluctuations (tick size), and the natural smallest quantity which can be traded (minimum lot size).

In physics, dimensional analysis begins with fundamental units of mass, distance, and time. In finance, dimensional analysis begins with fundamental units of time, currency, and shares (or contracts). In physics, dimensional analysis is often augmented by a conservation law based on principles of physics. In finance, proceeding further requires introducing conservation laws based on principles of finance. These conservation laws naturally take a form of no-arbitrage restrictions. The Black-Scholes option-pricing model[5] is based on the no-arbitrage principle that that all riskless trading strategies must earn the same riskfree return. Here we use a no-arbitrage principle closely related to the Modigliani-Miller theorem,[6] which states that changes in a firm's mix of equity and debt securities do not affect the economic outcomes associated with bets which transfer the risks embedded in the firm's securities. In other words, if a risky asset is combined with positive or negative amounts of an infinitely liquid riskless cash-equivalent asset and the bundle traded as a single package, then the economics behind these "package" trades would not depend on how much riskless asset is included in the package.

In the area of market microstructure, dimensional analysis leads to new insights which are neither obvious nor well-known. As discussed next, dimensional analysis makes it possible to describe various microscopic properties of financial markets in simple macroscopic terms.

1 Dimensional Analysis and Leverage Neutrality.

Trading is costly; bets tend to move market prices relative to pre-trade benchmarks. Buy bets push prices up and sell bets push prices down relative to pre-trade price levels. These adverse price movements, called "market impact," occur as a result of adverse selection; traders on the opposite side of the bet believe correctly that bets contain private information, and they compensate for this by requiring a price concession. Transaction cost models quantify trading costs. Good transaction cost models are of great interest to traders.

Suppose that the market impact cost of executing a bet of Q shares is a function of the number of shares Q, the stock price P, share volume V, returns variance σ^2, and dollar bet cost C. While P, V, and σ^2 are potentially observable, bet cost C may be difficult to observe empirically. For now, think of bet cost C as the unconditional expected dollar price impact cost of executing a bet of random size $\tilde{Q}$, with $\tilde{Q} > 0$ for buy bets and $\tilde{Q} < 0$ for sell bets. Thus, C measures how much a trader must pay for executing a random bet relative to a pre-trade benchmark.

Let $G := g(Q, P, V, \sigma^2, C)$ denote the price impact cost as a fraction of the value traded $P \cdot |Q|$; the quantity G is dimensionless, with $G \geq 0$. The quantities Q, P, V, σ^2, and C are measured using units of currency, shares, and time. If units of currency are reduced by a factor U, shares by a factor S, and time by a factor T, then $|Q|$ increases by a factor S, P increases by a factor US^{-1}, V increases by a factor ST^{-1}, σ^2 increases by a factor T^{-1}, and C increases by a factor U. Since there are three sets of distinct units and five dimensional quantities—Q, P, V, σ^2, C—it is possible to form two independent dimensionless quantities. Without loss of generality, let L and Z denote these dimensional quantities, defined by

$$L := \left(\frac{\theta \cdot P \cdot V}{\sigma^2 \cdot C} \right)^{1/3}, \qquad Z := \frac{P \cdot Q}{L \cdot C}. \tag{1}$$

Here θ is a dimensionless scaling constant, and the exponent of one-third in the definition of L is chosen strategically for important reasons related to leverage neutrality discussed below.

The variables L and Z have an intuitive interpretation. Without loss of generality, choose the scaling constant θ such that $E\{|\tilde{Z}|\} = 1$. Then $\tilde{Z}$ can be interpreted as "scaled bet size" because it expresses the size of a bet Q as a multiple of mean unsigned bet size $E\{|\tilde{Q}|\}$. The definition of $\tilde{Z}$ also implies $1/L = C/(E\{P \cdot |\tilde{Q}|\})$. Since the numerator C is the expected dollar cost of a bet and the denominator $E\{P \cdot |\tilde{Q}|\}$ is the expected dollar value of the bet, the variable $1/L$ measures the value-weighted expected market impact cost of a bet, expressed as a fraction of the dollar value traded. Illiquidity is a synonym for transactions costs; it is therefore reasonable to interpret $1/L$ as "illiquidity" and L as "liquidity;" Greater liquidity is associated with larger bets since $E\{P \cdot |\tilde{Q}|\} = C \cdot L$.

Without loss of generality, re-define the arguments of the function g so that it is written as $g(P, Q, \sigma^2, L, Z)$. The arguments P, Q, σ^2 are dimensional quantities which trivially span the three units of currency, shares, and time. The two arguments L and Z are independent dimensionless quantities from which V and C can be recovered. Since the value of $g(P, Q, \sigma^2, L, Z)$ is itself dimensionless, it cannot depend on the dimensional quantities P, Q, and σ^2. Thus, *dimensional analysis* implies that the function g can be further simplified by writing it as $g(L, Z)$. Dimensional analysis alone leads to the result that the market impact cost of a bet of Q shares, expressed as a fraction of the bet value $P \cdot |Q|$, must be a function of the two variables L and Z, assuming the five arguments Q, P, V, σ^2, C define a correct specification for the function.

To refine the transaction cost model further, introduce a conservation law in the form of *leverage neutrality*. This is closely related to Modigliani-Miller equivalence, the idea that a change in leverage—the ratio of firm's assets to firm's equity—does not affect the underlying economics of the risk transfer represented by a bet of Q shares. Suppose that the stock is levered up by a factor A as a result of paying a cash dividend of $(1 - A^{-1}) \cdot P$ financed with cash or riskless debt. Since a bet of Q shares transfers the same economic risk, the number of shares in a bet Q does not change, and trading volume V does not change. Since the economic risk of a bet does not change, the dollar cost of the bet C does not change either. The ex-dividend price of a share is $A^{-1} \cdot P$ because the value of the share-plus-dividend is conserved. Each share continues to have the same dollar risk $P \cdot \sigma$; therefore, the returns standard deviation σ increases to $A \cdot \sigma$, and the returns variance σ^2 increases to $A^2 \cdot \sigma^2$. It is straightforward to verify that L changes to $A^{-1} \cdot L$ and Z remains unchanged. Strategically incorporating the exponent $1/3$ into the definition of L has the effect of making L scale inversely proportionally with A, just like P. The percentage cost G of executing a bet of Q shares changes by a factor A because the dollar cost of executing this bet remains unchanged while the dollar value of the bet itself declines from $P \cdot |Q|$ to $A^{-1} \cdot P \cdot |Q|$.

Leverage neutrality thus implies that for any A, the function g satisfies the homogeneity condition $g(A^{-1} \cdot L, Z) = A \cdot g(L, Z)$. Letting $A = L$, the function g can be written $g(L, Z) =$

$L^{-1} \cdot g(1, Z)$. Define the univariate function f by $f(Z) := g(1, Z)$. Now G can be written in the form $G = L^{-1} \cdot f(Z)$.

The dimensionless and leverage neutral combination $G \cdot L$ must be a function of the dimensionless and leverage neutral argument Z. In a more explicit form, we have

$$G = \frac{1}{L} \cdot f\left(\frac{P \cdot Q}{C \cdot L}\right). \tag{2}$$

or, in terms of the original five parameters,

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

Equation (3) describes a general specification for transaction costs functions consistent with the scaling implied by dimensional analysis and leverage neutrality. This general specification is consistent with different assumptions about the shape of the function f.

To summarize, the combination of dimensional analysis and leverage neutrality reduces the problem from determining the structure of a function of five parameters $g(Q, P, V, \sigma^2, C)$ to determining the structure of a function of only one parameter $f(Z)$. The percentage transaction costs (2) can be presented as the product of a dimensionless security-specific measure of illiquidity $1/L$ and a dimensionless function $f(Z)$ of scaled bet size Z, which is dimensionless and leverage neutral.

2 Market Microstructure Invariance.

Dimensional analysis does not generate operational market microstructure predictions per se. To obtain useful empirical predictions based on transaction costs model (3), it is necessary to think about how to measure relevant quantities. The derivation above refers to at least five quantities: asset price P, trading volume V, returns volatility σ, bet size Q, bet cost C, and possibly other measures of transactions costs such as bid-ask spreads. Three of the quantities—asset price P, trading volume V, and returns volatility σ—can be observed directly or readily estimated from public data feeds on securities transactions; these are observable characteristics of an asset. The size Q is a characteristic of a bet privately known to a trader. While bid-ask spreads can be observed from public data feeds, other estimates of transactions costs generally requires having confidential data on transactions which allows transactions of one trader to be distinguished from transactions of another. More ambiguous is the issue of how the cost of a bet C varies across assets.

Market microstructure invariance[7] is the empirical hypothesis that the dollar value C is the same for all time periods and for all assets such as stocks, bonds, commodities, foreign exchange, and derivatives. The a priori justification for this invariance hypothesis is Ockham's razor: it is is simplest possible hypothesis. This invariance hypothesis is neither an implication of dimensional analysis nor an implication of leverage neutrality. Instead, it is an economic hypothesis motivated by the intuition that asset managers allocate scarce intellectual resources across assets and across time in such a manner that the cost of bets is equated.[8] To apply invariance across markets with different currencies and real exchange rates, it is necessary to further generalize this hypothesis and, for example, scale the expected dollar cost C by the productivity-adjusted wages of finance professionals in the local currency to make this variable dimensionless.

In both physics and market microstructure, application of invariance principles requires that certain assumptions be met. For example, the laws of physics hold in simplest form for objects traveling in a vacuum but have to be modified when resistance from air generates friction. Similarly, in market microstructure, the invariance assumption may hold only under idealized conditions. For example, the predictions of invariance may hold most closely when tick size is small, market makers are competitive, and transactions fees and taxes are minimal. Invariance principles provide a benchmark from which the importance of frictions such as a large tick size, non-competitive market access, or high fees and taxes can be measured.

Under the invariance assumption, instead of having different models for different securities and different time periods, it is necessary to calibrate only one parameter C for all assets, not a different value of C for each asset. Together with calibration of the scaling constant θ and the shape of the invariant cost function $f(Z)$ in equation (2), the knowledge of the parameter C makes it possible to write an operational transaction costs model for any market. The constant C helps to relate the microscopic details of trading in a security to its macroscopic properties. For example, it helps to relate the microscopic transaction cost of a bet to observable volume V, price P, and volatility σ. Preliminary calibration based on portfolio transitions data suggests that cost C of a bet is approximately equal to 2,000 dollars.[7] This estimate is obtained from analysis based on a large sample of portfolio transitions orders; a portfolio transition occurs when a large investor, such as a pension plan sponsor, hires a professional third party to make the trades necessary to move assets from from one asset manager to another.

If invariance holds, the dimensionless liquidity index $L \sim (P \cdot V \cdot \sigma^{-2})^{1/3}$ is a natural, simple measure of liquidity which is easy to calculate using data on volume and volatility. This security-specific metric does not change when a stock splits or the frequency with which data is sampled changes.

The general specification for a transaction cost function (3) is consistent with different functional forms. Suppose that f is a power function of the form $f(Z) = \bar{\lambda} \cdot |Z|^{\omega}$. A proportional bid-ask spread cost ($\omega = 0$) implies

$$G = \text{const} \cdot \frac{1}{L}. \tag{4}$$

A linear market impact cost ($\omega = 1$), which is often assumed in theoretical models,[9] implies

$$G = \text{const} \cdot \frac{P \cdot |Q|}{L^2}. \tag{5}$$

A square-root market impact cost ($\omega = 1/2$) implies

$$G = \text{const } \cdot \sigma \cdot \left(\frac{|Q|}{V} \right)^{1/2}. \tag{6}$$

Empirical estimates often support the square root specification with a proportionality factor close to one.[10]. Equation (3) is a general structural transaction costs model. Equations (4), (5), and (6) are special cases consistent with invariance.

3 Empirical Evidence based on Bid-Ask Spread in the Russian equities market.

Dimensional analysis and leverage neutrality imply a scaling law for the quoted bid-ask spread, which measures the difference between the highest price at which a trader is willing to buy ("bid") and the lowest price at which a trader is willing to sell ("offer"). Let S denote the bid-ask spread measured in the same units as price P. From equation (4), market microstructure invariance implies

$$\log\left(\frac{S}{P}\right) = \text{const} + 1 \cdot \log(1/L). \tag{7}$$

For empirical estimation, the unknown invariant constants C and θ can be factored out of the definition of L and incorporated into the constant term in equation (7). The coefficient of one on $\log(1/L)$ implies a scaling exponent of $-1/3$ on $P \cdot V \cdot \sigma^{-2}$.

To test this relationship, we use data from the Moscow Exchange for January–December 2015 provided by Interfax Ltd. The data cover 50 Russian stocks in the RTS index as of June 15, 2015. The five largest companies are Gazprom, Rosneft, Lukoil, Novatek, and Sberbank. The Russian stock market is centralized with all trading implemented in a consolidated limit-order book. Since the tick size was small during this period, this market friction was less likely to affect the bid-ask spreads of Russian stocks than in other markets with larger tick size, like the U.S. The lot size was small in the Russian stock market as well. For each of the 50 stocks and each of the 250 trading days, the average percentage spread is calculated as the mean of the percentage spread at the end of each minute during trading hours from 10:00 to 18:50. The realized volatility is calculated based on summing squared one-minute changes in the mid-point between the best bid and best offer prices at the end of each minute during trading hours. Table 1 presents summary statistics for this sample.

Figure 1 plots the log bid-ask spread $\log(S/P)$ against $\log(1/L)$. Each of 12,426 points represents the average bid-ask spread for one stock for one day. Different colors represent different stocks. For comparison, we add a solid line $\log(S/P) = 2.112 + 1 \cdot \log(1/L)$, where the slope is fixed at level predicted by market microstructure invariance and the intercept is estimated. All observations cluster around this benchmark line.

On aggregate sample, the fitted line is $\log(S/P) = 2.093 + 0.998 \cdot \log(1/L)$, with standard errors of estimates 0.040 and 0.005, respectively; the R-square is 0.876. The invariance prediction that the slope coefficient is one is not statistically rejected. The fitted line for a similar regression over monthly averages instead of daily averages is $\log(S/P) = 2.817 + 1.078 \cdot \log(1/L)$ with standard errors of estimates 0.164 and 0.019, respectively; its R-square is 0.923. The invariance prediction that the slope coefficient is one is statistically rejected in this case, but remains economically close to the data.

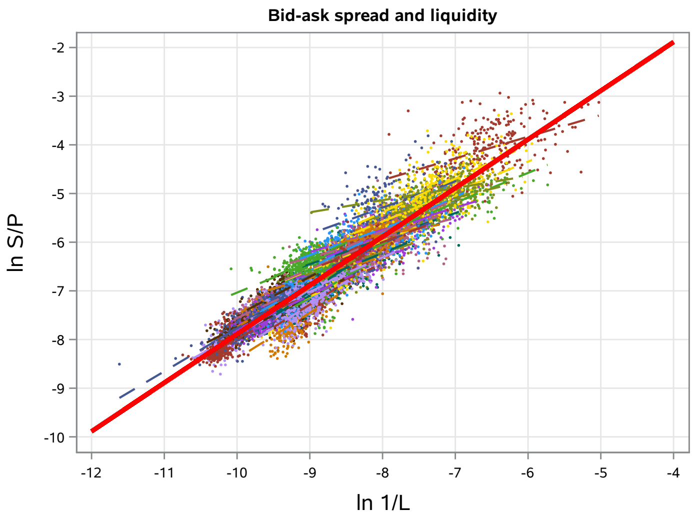

Figure 1: Bid-ask spread and liquidity.

The 50 dashed lines in figure 1 are fitted based on data for the 50 individual securities. The slopes, which vary from 0.249 to 1.011, are substantially lower than the invariance-implied slope of one, which is indistinguishable from the fitted line for the aggregate data. A possible explanation is that a substantial part of the variation in stock-specific measures of trading activity is due to variations in trading activity of the overall market and may therefore be only loosely related to variations in bid-ask spreads. The downward bias is more pronounced for less liquid stocks, suggesting that it may also be related to correlation between explanatory variables and error terms due to natural endogeneity in contemporaneous variables.

4 Additional Applications and Extensions.

The empirical implications of dimensional analysis, leverage invariance, and market microstructure invariance can be generalized. The analysis above assumes that function g is correctly specified in terms of only five parameters Q, P, V, σ^2, C. The approach can be easily extended to include additional function arguments. It can also be extended to derive scaling laws for other variables. Here are several illustrations.

The magnitude of transaction costs is likely to depend on other variables. Traders believe that order execution costs are lower if the execution horizon is longer. Transaction costs may also depend on market frictions such as tick size and minimum round lot size. The tick size for U.S. stocks is generally one cent, and the minimum round lot size is generally 100 shares; for Russian stocks there is more variation in these parameters.

First, add to the original five parameters the horizon of execution H measured in units of time, the tick size K_{MIN} measured in dollars per share, and the minimum round lot size Q_{MIN} measured in shares. Second, re-scale H, K_{MIN}, and Q_{MIN} to make them dimensionless and leverage neutral. There is a unique way to do this using the four variables P, V, σ^2, and C (including the liquidity variable L). The unique, re-scaled values are $H \cdot V/|Q|$, $K_{MIN} \cdot L/P$, and $Q_{MIN} \cdot \sigma^2 \cdot L^2/V$, respectively (up to constants of proportionality). Equation (2) then becomes

$$G = \frac{1}{L} \cdot f\left(\frac{P \cdot Q}{C \cdot L}, \frac{H \cdot V}{|Q|}, \frac{K_{MIN} \cdot L}{P}, \frac{Q_{MIN} \cdot \sigma^2 \cdot L^2}{V}\right). \tag{8}$$

This more general specification remains consistent with scaling laws but allows for non-linear relationships among the different arguments of f. Other variables can be added to the transaction costs model analogously.

Optimal execution horizon is of obvious interest to traders. Suppose that the optimal (cost-minimizing) execution horizon H^* for an order of Q shares depends on P, V, σ^2, C, K_{MIN}, and Q_{MIN}. Since the ratio $H^* \cdot V/|Q|$ is dimensionless and leverage neutral, the same logic as above implies the following formula

$$H^* = \frac{|Q|}{V} \cdot h^*\left(\frac{P \cdot Q}{C \cdot L}, \frac{K_{MIN} \cdot L}{P}, \frac{Q_{MIN} \cdot \sigma^2 \cdot L^2}{V}\right). \tag{9}$$

When tick size is large, larger quantities available at the best bid and offer may make the execution horizon shorter. If tick size and minimum lot size do not affect execution horizon, this horizon H^* depends only on scaled bet size $Z := P \cdot Q/(C \cdot L)$. If the function h^* is a constant, then it is optimal to choose the execution horizon so that traders execute all trades as the same fraction of volume, say one percent of volume until execution of the bet is completed.

Setting optimal tick size and minimum lot size is of interest for exchange officials and regulators. Since the scaled quantities $K^*_{MIN} \cdot L/P$ and $Q^*_{MIN} \cdot L^2 \cdot \sigma^2/V$ are dimensionless and leverage neutral, the scaling laws for these market frictions can be written as

$$K^*_{MIN} = \text{const} \cdot \frac{P}{L}, \qquad Q^*_{MIN} = \text{const} \cdot \frac{V}{L^2 \cdot \sigma^2}. \tag{10}$$

Since the proportionality constant does not vary across securities, these measures provide good benchmarks for comparing the restrictiveness of actual tick size and minimum lot size across securities.

If exchanges set tick size and minimum lot size at their optimal levels of K^*_{MIN} and Q^*_{MIN} and traders choose optimal execution horizons H^*, then f in equation (8) becomes a function of only one argument Z again.

Our approach can be also used to derive more general scaling laws for the bid-ask spread. The bid-ask spread is an integer number of ticks which fluctuates as trading occurs. Let S denote the average bid-ask spread, measured in dollars per share. Assume the average spread depends on P, V, σ^2, C, K_{MIN}, and Q_{MIN}. Dimensional analysis and leverage neutrality imply that the re-scaled spread $S \cdot L/P$, which is dimensionless and leverage neutral, is a function s of only the re-scaled dimensionless and leverage-neutral variables K_{MIN} and Q_{MIN}:

$$\frac{S}{P} = \frac{1}{L} \cdot s\left(\frac{K_{MIN} \cdot L}{P}, \frac{Q_{MIN} \cdot \sigma^2 \cdot L^2}{V}\right). \tag{11}$$

If tick size and minimum lot size have no influence on quoted bid-ask spreads, then the the relationship simplifies to $S/P \sim 1/L$, as tested above for the Russian equities market.

Our approach can be also used to derive more general scaling laws for trading data. A bet of size Q may be executed as a large number of smaller trades. Let Q_T denote a trade, a fraction of a bet. Trades and bets have the same units but different economics. While it is reasonable to conjecture that the size of bets does not depend on tick size or minimum lot size, the size of trades into which bets are "shredded" will obviously depend on both of them. With large tick size, there will typically be large quantities available at the bid and offer; therefore, large bets may be executed as trades of large size which clean out available bids and offers. Empirical evidence suggests that trades have become so small in recent years that minimum lot size is often a binding constraint.[11] Since $\tilde{Z} := P \cdot \tilde{Q}_T/(C \cdot L)$ is dimensionless and leverage-neutral, the same analysis as above leads to the following scaling laws for the probability distribution of $\tilde{Q}_T$:

$$\text{Prob}\left\{\frac{P \cdot \tilde{Q}_T}{C \cdot L} < Z\right\} = F_Q\left(Z, \frac{K_{MIN} \cdot L}{P}, \frac{Q_{MIN} \cdot \sigma^2 \cdot L^2}{V}\right). \tag{12}$$

Similar scaling laws can be derived for the quantities at the best bid and offer—$\tilde{Q}_B$ and $\tilde{Q}_A$—as well as for the entire limit order book.

Let γ_T denote the number of transactions per day. After re-scaling, $\gamma_T/(\sigma^2 \cdot L^2)$ becomes dimensionless and leverage-neutral, further suggesting a scaling law of the form

$$\gamma_T = \sigma^2 \cdot L^2 \cdot f\left(\frac{K_{MIN} \cdot L}{P}, \frac{Q_{MIN} \cdot \sigma^2 \cdot L^2}{V}\right). \tag{13}$$

If tick size and minimum lot size do not affect the trading process, then the number of transactions satisfies $\gamma_T \sim \sigma^2 \cdot L^2$, and their average size satisfies $P \cdot E\{|\tilde{Q}_T|\} = C \cdot L$.

Figure 2 presents results of testing the prediction $\gamma_T \sim \sigma^2 \cdot L^2$ using the data from the Moscow Exchange. The figure has 12,426 points plotting the log number of transactions $\log(\gamma_T)$ against $\log(\sigma L)$ for each of 50 stocks and each of 250 days. A benchmark line $\log(\gamma_T) = -1.937 + 2 \cdot \log(\sigma L)$, where the slope is fixed at a predicted level of two and intercept is estimated, is added for comparison. The results for the aggregate sample are broadly consistent with the predicted slope of two. The fitted line is $\log(\gamma_T) = -3.085 + 2.239 \cdot \log(\sigma L)$ with standard errors of estimates equal to 0.038 and 0.008, respectively; its R-square is 0.882. As before, the slopes of fitted lines for individual stocks are systematically lower, ranging from 1.156 to 1.795 and depicted with dashed lines.

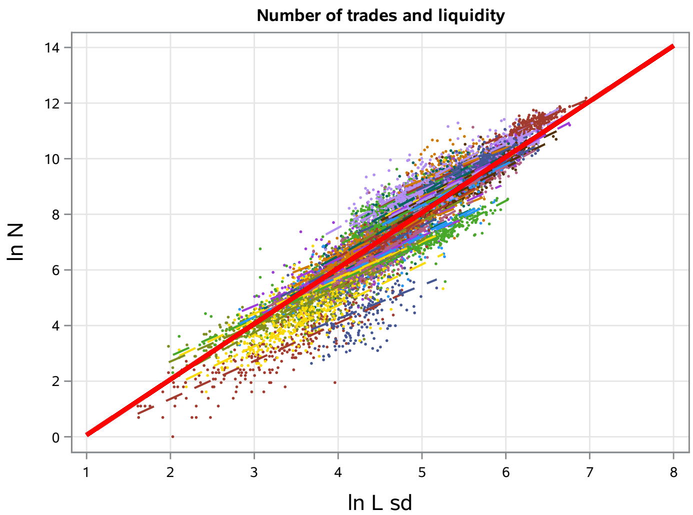

Figure 2: Number of trades per day and liquidity.

There is a growing empirical evidence that the scaling laws discussed above match patterns in financial data, at least approximately. These scaling laws are found in data on transaction costs and order size distributions for institutional orders,[7] in data on trades executed in the U.S. and South Korean equities markets,[11,12] in data on news articles published by Thomson Reuters,[13] and in intraday trading patterns of the S&P E-mini futures market.[14]

Checking the validity of invariance predictions in other samples, improving the accuracy of estimates, and the triangulation of proportionality constants are important tasks for future research.

Our research here is relevant for risk managers and traders, who seek to minimize and measure market impact costs. It also establishes politically neutral, scientific benchmarks for numerous policy issues connected with market microstructure such as setting tick sizes and minimum lot sizes as well as position limits, margin requirements, and repo haircuts. Such research is highly relevant for the economic analysis of market crashes,[15] such as the U.S. stock market "flash crash" of May 2010[16] or the U.S. bond market "flash rally" of October 2014.[17] Lastly, it directly relates to designing liquidity management tools, one of the central issues addressed by the Dodd-Frank Act and Basel III regulatory initiatives.

1. Bouchaud, J.-P., Farmer, D. J. & Lillo, F. How markets slowly digest changes in supply and demand. In *Handbook of Financial Markets: Dynamics and Evolution*, 47–160 (North Holland, 2009).

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

3. Barenblatt, G. I. *Scaling, Self-Similarity, and Intermediate Asymptotics* (Cambridge University Press, 1996).

4. Kolmogorov, A. N. On the log-normal distribution of particle sizes during break-up process. *Doklady Akademii Nauk SSSR* **31**, 99–101 (1941b).

5. Black, F. & Scholes, M. The pricing of options and corporate liabilities. *The Journal of Political Economy* 637–654 (1973).

6. Modigliani, F. & Miller, M. H. The cost of capital, corporation finance and the theory of investment. *American Economic Review* **48**, 261–297 (1958).

7. Kyle, A. S. & Obizhaeva, A. A. Market Microstructure Invariance: Empirical Hypotheses. *Econometrica* accepted for publication (2016).

8. Kyle, A. S. & Obizhaeva, A. A. Market Microstructure Invariance: A Dynamic Equilibrium Model. Tech. Rep., Working Paper, University of Maryland (2016).

9. Kyle, A. S. Continuous auctions and insider trading. *Econometrica* **53**, 1315–1335.

10. Grinold, R. C. & Kahn, R. N. In *Active Portfolio Management* (Probus Publishing, 1995).

11. Kyle, A. S., Obizhaeva, A. A. & Tuzun, T. Microstructure Invariance in U.S. Stock Market Trades. Tech. Rep., Working Paper, University of Maryland (2016).

12. Bae, K., Kyle, A. S., Lee, E. J. & Obizhaeva, A. A. An Invariance Relationship in the Number of Buy-Sell Switching Points. Tech. Rep., Working Paper, University of Maryland (2014).

13. Kyle, A. S., Obizhaeva, A. A., Sinha, N. R. & Tuzun, T. News Articles and the Invariance Hypothesis. Tech. Rep., Working Paper, SSRN-id1786124 (2014).

14. Andersen, T. G., Bondarenko, O., Kyle, A. S. & Obizhaeva, A. A. Intraday Trading Invariance in the E-mini S&P 500 Futures Market. Tech. Rep., Working Paper, SSRN-id2693810 (2016).

15. Kyle, A. S. & Obizhaeva, A. A. Large Bets and Stock Market Crashes. Tech. Rep., Working Paper, University of Maryland (2016).

16. Staffs of the CFTC and SEC. *Findings Regarding the Market Events of May 6, 2010* (2010b). Report of the Staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues. September 30, 2010. Final Report.

17. Staffs of the Fed, the CFTC, and SEC. *The U.S. Treasury Market on October 15, 2014* (2015). Joint staff report by U.S. Department of the Treasury, Board of Governors of the Federal Reserve System, Federal Reserve Bank of New York, U.S. Securities and Exchange Commission, U.S. Commodity Futures Trading Commission. July 13, 2015.

	avg	p5	p50	p95	p100
Cap (rub)	476	24	189	1,904	3,420
Cap (usd)	8.77	0.44	3.50	35.02	63.02
$V \cdot P$ (rub)	542	3	73	2,607	6,440
σ	0.019	0.014	0.018	0.029	0.032
S/P	19	3	12	61	129
γ_T	7,328	65	2,792	22,169	71,960

Table 1: The table presents summary statistics (average values and percentiles) for the sample of 50 Russian stocks: ruble and dollar capitalization Cap (in billions), average daily volume $V \cdot P$ in millions of rubles, daily returns volatility σ, average percentage spread S/P in basis points (hundredths of a percent), and average number of trades per day γ_T as of June 2015.