



Are trading invariants really invariant? Trading costs matter

Frédéric Bucci, Fabrizio Lillo, Jean-Philippe Bouchaud, Michael Benzaquen

► To cite this version:

Frédéric Bucci, Fabrizio Lillo, Jean-Philippe Bouchaud, Michael Benzaquen. Are trading invariants really invariant? Trading costs matter. Quantitative Finance, 2020. hal-02323318

HAL Id: hal-02323318

<https://hal.science/hal-02323318>

Submitted on 21 Oct 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Are trading invariants really invariant? Trading costs matter

Frédéric Bucci ^{*1,7}, Fabrizio Lillo^{2,3}, Jean-Philippe Bouchaud^{4,5,7}, and Michael Benzaquen^{4,6,7}

¹Scuola Normale Superiore di Pisa, Piazza dei Cavalieri 7, 56126 Pisa, Italy

²Department of Mathematics, Università di Bologna, Piazza di Porta San Donato 5, 40126 Bologna, Italy

³CADS, Human Technopole, Milan, Italy

⁴Capital Fund Management, 23 rue de l'Université, 75007, Paris, France

⁵CFM-Imperial Institute of Quantitative Finance, Department of Mathematics, Imperial College, 180 Queen's Gate, London SW7 2RH

⁶Ladhyx UMR CNRS 7646 & Department of Economics, Ecole polytechnique, 91128 Palaiseau Cedex, France

⁷Chair of Econophysics and Complex Systems, Ecole polytechnique, 91128 Palaiseau Cedex, France

February 12, 2019

Abstract

We revisit the trading invariance hypothesis recently proposed by Kyle and Obizhaeva [1] by empirically investigating a large dataset of bets, or metaorders, provided by ANcerno. The hypothesis predicts that the quantity $I := \mathcal{R}/N^{3/2}$, where \mathcal{R} is the exchanged risk (volatility \times volume \times price) and N is the number of bets, is invariant. We find that the 3/2 scaling between \mathcal{R} and N works well and is robust against changes of year, market capitalisation and economic sector. However our analysis clearly shows that I is not invariant. We find a very high correlation $R^2 > 0.8$ between I and the total trading cost (spread and market impact) of the bet. We propose new invariants defined as a ratio of I and costs and find a large decrease in variance. We show that the small dispersion of the new invariants is mainly driven by (i) the scaling of the spread with the volatility per transaction, (ii) the near invariance of the distribution of metaorder size and of the volume and number fractions of bets across stocks.

Contents

1	Introduction	2
2	Data	3
3	The 3/2-law	3
3.1	Exchanged risk	3
3.2	Empirical evidence	5
4	The trading invariant	5
4.1	Trading costs and trading invariants	5
4.2	Origin of the small dispersion of the new invariants	6
5	Conclusions	7
A	Statistics of metaorder sample	10
B	The 3/2-law under the microscope	10
C	Statistics of trading costs	11

*Corresponding author: frederic.bucci@sns.it

1 Introduction

Finding universal scaling laws between trading variables is highly valuable to make progress in our understanding of financial markets and market microstructure. In the wake of these discoveries, Kyle and Obizhaeva posit a *trading invariance principle* that must be valid for a *bet*, theoretically defined as a sequence of orders with a fixed direction (buy or sell) belonging to a single trading idea [1, 2]. This principle supports the existence of a universal invariant quantity I – expressed in dollars, independent of the asset and constant over time – which represents the average cost of a single *bet*. In particular, taking the share price P (in dollars per share), the square daily volatility σ_d^2 (in %² per day), the total daily amount traded with *bets* V (shares per day) and the average volume of an individual *bet* Q (in shares) as relevant variables, dimensional analysis suggests a relation of the form:

$$\frac{PQ}{I} = f\left(\sigma_d^2 \frac{Q}{V}\right), \quad (1)$$

where f is a dimensionless function. Invoking the Modigliani-Miller capital structure irrelevance principle yields $f(x) \sim x^{-1/2}$, which implies up to a numerical factor that:

$$I = \frac{\sigma_d P Q^{3/2}}{V^{1/2}} := \frac{\mathcal{R}}{N^{3/2}}, \quad (2)$$

where $\mathcal{R} := \sigma_d P V$ measures the total dollar amount of risk traded per day (also referred to as total exchanged risk or trading activity) while $N := V/Q$ represents the number of daily *bets* for a given contract. Notwithstanding, the 3/2 law can be interpreted with different degrees of universality as discussed in [3]: no universality (the 3/2 holds for some contracts only), weak universality (the 3/2 holds but with a non-universal value of I) and strong universality (the 3/2 holds and I is constant across assets and time).¹

Let us stress that identifying an elementary *bet* in the market is not a straightforward task. Theoretically, a bet is defined as a trading idea typically executed in the market as many trades over several days. As suggested by Kyle and Obizhaeva in their original work [1], metaorders, i.e. a bundle of orders corresponding to a single trading decision typically traded incrementally through a sequence of child orders, can be considered a proxy of these *bets*; in this work we will make use of such an approximation and use the words ‘bet’ or ‘metaorder’ indifferently. Beyond the subtleties in the bet’s definition, there has been in the past few years empirical evidence that the scaling law discussed above matches patterns in financial data, at least approximately. The 3/2-law was empirically confirmed by Kyle and Obizhaeva using portfolio transition data related to rebalancing decisions made by institutional investors and executed by brokers [1]. Andersen *et al.* [4] reformulated suitably the *trading invariance hypothesis* at the single-trade level and showed that the equivalent version of Eq. (2) in such a setting holds remarkably well using public trade-by-trade data relative to the E-mini S&P 500 futures contracts. Benzaquen *et al.* [3] substantially extended these empirical results showing that the 3/2-law holds very precisely across 12 futures contracts and 300 single US stocks, and across a wide range of time scales. Recently, Pohl *et al.* [5] provided additional empirical evidence that the intriguing 3/2-law holds on trades data from the NASDAQ stock exchange.

Notwithstanding, empirical data at the single transaction scale – see in particular [3] – revealed that while the 3/2 law is very robust, the invariant I is actually quite far from invariant, as it varies from one asset to the other and across time, thus in favour of the *weak universality* degree. Note that this is consistent with the idea that a universal invariant with dollar units would be quite strange, given that the value of the dollar is itself time-dependent. Benzaquen *et al.* [3] showed that a more suitable candidate for an invariant was actually the dimensionless $\mathcal{I} := I/C$ where C denotes the spread trading costs.

Yet, single transactions are typically not the same as single bets. Large and medium sized orders are typically split in multiple transactions and traded incrementally over long periods of time. Empirically market data do not allow to infer the trading decision and to link different

¹ Note that here we only explore the daily level, time does not mean the same thing as in [3] where we varied the time intervals over which the variables were computed.

transactions to a single execution.² In order to test the trading invariance hypothesis at the bet level and its relation with trading costs, it is necessary to have a dataset of market-wide (i.e. not from a single institution) metaorders.

This is precisely the aim of the present paper, which leverages on a heterogeneous dataset of metaorders extracted from the ANcerno database.³ To our knowledge such a thorough analysis at the *bet* level for a wide range of assets is still lacking.

Our main finding is that, while the 3/2-law works surprisingly well, the quantity I is not invariant, as pointed out in [3]. We show that this quantity is strongly correlated with transaction costs, including spread and impact. Therefore we introduce new invariants, obtained by dividing I by the cost and we show that these quantities fluctuate very little across stocks and time periods. Finally we show that the observed small dispersion of the new invariants can be connected with three microstructural properties: (i) the linear relation between spread and volatility per transaction; (ii) the near invariance of the metaorder size distribution, and (iii) of the total volume and number fractions of the bets across different stocks.

The paper is organized as follows. In section 2 we describe the dataset collecting trading decisions of institutional investors operating in the US equity market. In section 3 we show that the 3/2-law holds surprisingly well at the daily level independently of the time period, of the market capitalisation and of the economic sector. In section 4 we compute the invariant I and we argue in favour of *weak universality*. We propose a more natural definition for a trading invariant that accounts both for the spread and the market impact costs; and we exhibit the microstructural origin of its small dispersion. Some conclusions and open questions are presented in section 5.

2 Data

Our analysis relies on a database made available by ANcerno, a leading transaction-cost analysis provider (www.ancerno.com). Our dataset counts heterogeneous institutional investors placing large buy or sell orders executed by a broker as a succession of smaller orders belonging to the same trading decision of a single investor. Our sample includes the period January 2007 – June 2010 for a total of 880 trading days. Only metaorders completed within at most a single trading day are held. Further, we select stocks belonging to the Russell 3000 index, thereby retaining ~ 8 million metaorders distributed quite uniformly in time and representing $\sim 5\%$ of the total reported market volume, regardless of market capitalisation (large, mid and small) and economical sectors (basic materials, communications, consumer cyclical and non-cyclical, energy, financial, industrial, technology and utilities). More details and statistics on the investigated sample are presented in Appendix A.

3 The 3/2-law

Here we investigate the trading invariance hypothesis at the daily level. The daily timescale choice avoids an elaborate analysis of when precisely each metaorder starts and ends, thereby averaging out all the non-trivial problems related to the daily simultaneous metaorders executed on the same asset [6].

3.1 Exchanged risk

From the metaorders executed on the same stock during the same day we compute the total exchanged volume in dollars: $\sum_{i=1}^N p_i v_i$, where N is the number of daily metaorders per asset in the ANcerno database, v_i and p_i are respectively the number of shares and the volume weighted

²In fact, for example, Kyle and Obizhaeva tackled this problem investigating a proprietary dataset of portfolio transitions.

³ANcerno Ltd. (formerly the Abel Noser Corporation) is a widely recognised consulting firm that works with institutional investors to monitor their equity trading costs. Its clients include many pension funds and asset managers. In [1] the authors claim that the ANcerno database includes more orders than the data set of portfolio transitions they used in their work. From a preliminary research Albert S. Kyle and Kingsley Fong found that proxies for *bets* in ANcerno data have size patterns consistent with the proposed invariance hypothesis discussed in [1, 2].

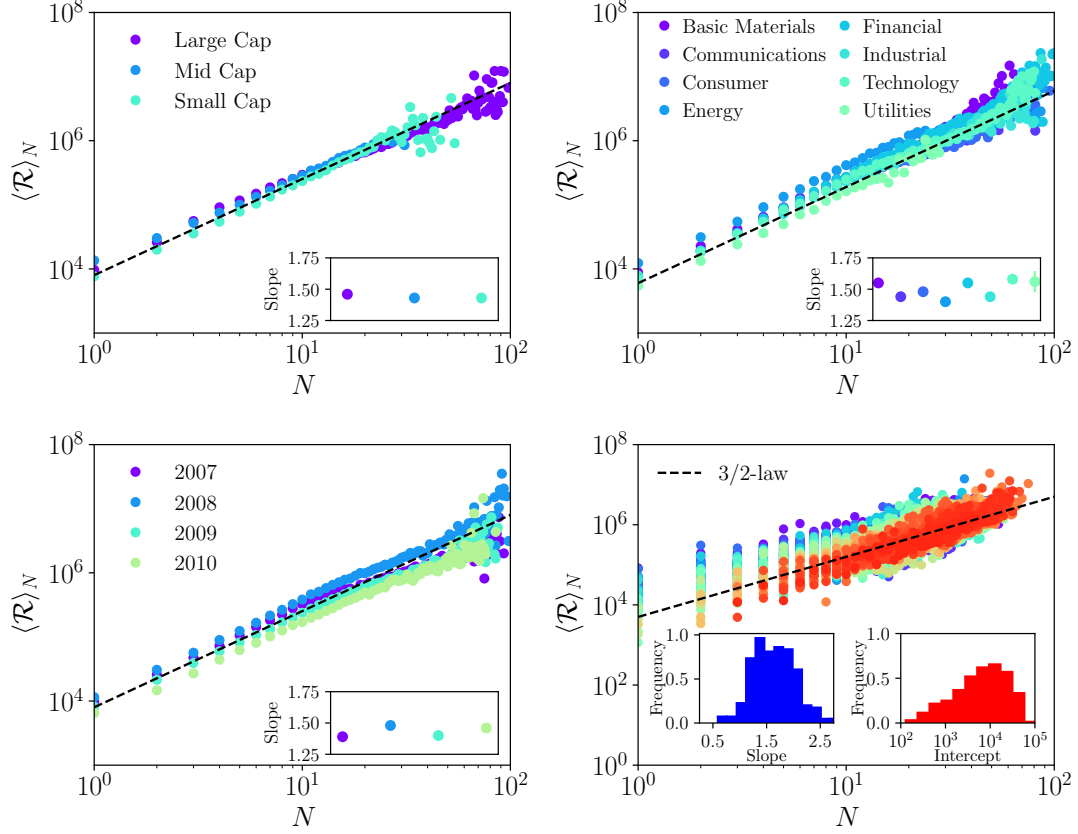


Figure 1: Plots of the mean daily exchanged risk $\langle \mathcal{R} \rangle_N$ as function of the daily number N of metaorders per asset conditional to the market capitalisation (top left panel), the economic sector (top right panel), and the time period (bottom left panel). The insets show the slopes obtained from linear regression of the data, firstly averaged respect to N and secondly log-transformed. The bottom right panel shows a plot of $\langle \mathcal{R} \rangle_N$ as function of N for a subset of 200 stocks chosen randomly from the pool of around three thousand US stocks: the two insets represent respectively the distribution of the slopes and of the y -intercept, i.e. $\langle I \rangle := \langle \mathcal{R} \rangle_N / N^{3/2}$, obtained from linear regression of the data, firstly averaged respect to N and secondly log-transformed of the data considering each stock separately.

average price (vwap) of the i -th available metaorder. We then define the total daily exchanged ANcerno risk per asset as:

$$\mathcal{R} := \sum_{i=1}^N \mathcal{R}_i, \quad \text{with} \quad \mathcal{R}_i = \sigma_d v_i p_i, \quad (3)$$

and where σ_d denotes the daily volatility per asset, computed as $\sigma_d = (p_{\text{high}} - p_{\text{low}}) / p_{\text{open}}$ from the high, low, and open daily prices only.⁴ The statistical properties of the bets, in terms of their associated risk \mathcal{R}_i and of their total daily number N per asset are discussed in Appendix A. The variability of the observables over several orders of magnitude should allow to test the 3/2-law quite convincingly.

⁴We checked that the results discussed in the present work are still valid using other definitions of the daily volatility and of the price in analogy to what done for example in [1]. Specifically, the results are still valid when computing σ_d with the Rogers-Satchell volatility estimator [3, 7] or as the monthly averaged daily volatility, i.e. $\bar{\sigma}_d = \sum_{m=1}^{25} \sigma_{d,m}$ and/or defining the price p_i as the closing price of the day before the metaorder's execution.

3.2 Empirical evidence

We introduce the mean daily exchanged risk $\langle \mathcal{R} \rangle_N$, where in general $\langle \bullet \rangle_N := \mathbb{E}[\bullet|N]$ denotes in a compact way the average over various days and stocks with a fixed daily number of metaorders N . As shown in the first three panels of Fig. 1 the scaling $\langle \mathcal{R} \rangle_N \sim N^{3/2}$ holds well independently of the conditioning to market capitalisation, economic sector, and time period. Slight deviations may have different origins but can mostly be attributed to the heterogeneous sample's composition in terms of stocks for each bucket in N . The 3/2-law is also valid for individual stocks, as shown in the bottom right panel of Fig. 1. The bottom left inset shows that the slopes are clustered around the 3/2-value. However, the y -intercepts of the fitted lines for individual stocks vary substantially (see the bottom right inset in the bottom right panel of Fig. 1), which indicates that I is not constant across different stocks. More empirical insights on the origin of the 3/2 law are presented in Appendix B.

4 The trading invariant

The conjecture that the quantity $\langle I \rangle := \langle \mathcal{R} \rangle_N / N^{3/2}$ is invariant across different contracts is clearly rejected by the empirical analysis performed in the previous section. Indeed, the quantity $\langle I \rangle$ varies by at least one order of magnitude across different stocks. This result goes against the *strong universality* version of the trading invariance hypothesis which states that both the average value $\langle I \rangle$ and the full probability distribution of $I = \mathcal{R} / N^{3/2}$ should be invariant across products. Dimensionally I is a cost (i.e. it is measured in dollars) and indeed the trading invariance hypothesis posits that the cost of a bet is invariant. Using the identification of metaorders and bets we can use the ANcerno dataset to estimate the trading cost, including a spread and a market impact component. We will show that I and trading cost are very correlated, and therefore propose new invariants based on their ratio.

4.1 Trading costs and trading invariants

Trading costs are typically divided into fees/commissions, spread, and market impact. For large orders, like those investigated here, fees/commissions typically account for a very small fraction and therefore we will neglect them. We shall however take into consideration both the spread cost (as was done at the single-trade level in [3]) and the market impact cost computed from the square root law (see *e.g.* [6, 8, 9, 10, 11, 12, 13, 14]). We thus define the average daily *bet's* trading cost




Figure 2: (Left) Empirical distributions of the KO invariant $I = \mathcal{R}N^{-3/2}$, of the daily average *bet's* total trading cost \mathcal{C} (using $Y_{\text{spd}} = 3.5$ and $Y_{\text{imp}} = 1.5$ in Eq. 4), and of the dimensionless invariant $\mathcal{I} := I/\mathcal{C}$. (Right) Empirical distributions in log-log scale of the KO invariant I rescaled respectively by the total daily average cost \mathcal{C} , by the spread cost \mathcal{C}_{spd} and by the market impact cost \mathcal{C}_{imp} .

as:

$$\mathcal{C} = \mathcal{C}_{\text{spd}} + \mathcal{C}_{\text{imp}} = Y_{\text{spd}} \times \frac{1}{N} \sum_{i=1}^N S v_i + Y_{\text{imp}} \times \frac{1}{N} \sum_{i=1}^N \sigma_d v_i p_i \sqrt{\frac{v_i}{V_d}} := Y_{\text{spd}} \times \mathcal{C}_{\text{spd}}^0 + Y_{\text{imp}} \times \mathcal{C}_{\text{imp}}^0, \quad (4)$$

with S the average daily spread,⁵ V_d the total daily market volume, Y_{spd} and Y_{imp} two constants to be determined. The factor Y_{spd} depends, among other things, on the fraction of trades of the metaorder executed with market orders, whereas Y_{imp} only weakly depends on the execution algorithm and is typically estimated to be very close to unity [6, 8, 15]. The empirical properties of $\mathcal{C}_{\text{spd}}^0$ and $\mathcal{C}_{\text{imp}}^0$ and the relative importance of the two terms as a function of the metaorder size are presented in Appendix C. To determine Y_{spd} and Y_{imp} we perform an ordinary least square regression of the KO invariant I with respect to the daily average cost \mathcal{C} defined for each asset by Eq. 4. We obtain $Y_{\text{spd}} \simeq 3.5$, $Y_{\text{imp}} \simeq 1.5$ and a coefficient of determination $R^2 \simeq 0.8$. These results show that the original KO invariant is indeed strongly correlated with the trading cost. Since these costs have no a priori reason to be universal, this explains why I is not invariant.

Guided by such results and by the fact that a market microstructure invariant, if any, should be dimensionless, we define new invariants by dividing the original KO invariant I by the cost of trading. Therefore, we consider three different specifications, namely:

$$\mathcal{I} = \frac{I}{\mathcal{C}}, \quad \mathcal{I}_{\text{spd}} = \frac{I}{\mathcal{C}_{\text{spd}}}, \quad \mathcal{I}_{\text{imp}} = \frac{I}{\mathcal{C}_{\text{imp}}}. \quad (5)$$

The left panel of Figure 2 shows the distribution of the original KO invariant I together with that of \mathcal{I} , and of the cost \mathcal{C} . It is visually quite clear that rescaling by the cost dramatically reduces the dispersion, and that the distribution of I is very similar to that of \mathcal{C} , with some deviation for small value. The right panel compares the distribution of \mathcal{I} with that of the other two new invariants. A quantitative comparison is provided in Table 1, which reports the mean, the standard deviation, and the coefficient of variation⁶ (CV) of I and of the three new invariants. It is clear that, due to the correlation between I and \mathcal{C} , the new invariants \mathcal{I}_\circ (with $\circ = \text{spd}, \text{imp}$) have a much smaller CV than I . Since the distributions have clear fat tails, we also implemented a robust version of CV obtained by replacing the standard deviation with the mean absolute deviation (MAD), here denoted CV_{MAD} . The table indicates that also in this case the new invariants are much more peaked than I . Notice also that with CV_{MAD} the three new invariants become similar, while with CV the invariant \mathcal{I}_{imp} is more dispersed.

Table 1: Statistics of the different invariants, namely the original KO invariant I (left), and the three new ones rescaled by cost (right). MAD is the mean absolute deviation and CV stands for coefficient of variation.

	$I \cdot 10^3$ (\$)	\mathcal{I}	\mathcal{I}_{spd}	\mathcal{I}_{imp}
mean	6.33	2.20	4.70	7.8
st. dev.	11	1.84	3.11	12.2
MAD	6.9	1.25	2.21	7.56
CV	1.74	0.84	0.66	1.56
CV_{MAD}	1.09	0.57	0.47	0.97

4.2 Origin of the small dispersion of the new invariants

Here we investigate the origin of the small dispersion of the new invariants. Let us first consider only the impact cost normalisation only and rewrite \mathcal{I}_{imp} as:

$$\mathcal{I}_{\text{imp}} = \frac{N \sum_{i=1}^N \sigma_d p_i v_i}{Y_{\text{imp}} N^{3/2} (\sigma_d \sum_{i=1}^N p_i v_i \sqrt{v_i/V_d})}. \quad (6)$$

⁵The daily spread is recovered from a dataset provided by CFM since it is not available in the ANcerno dataset.

⁶The coefficient of variation is the ratio of standard deviation and mean, an indicator of distribution ‘peakedness’.




Figure 3: Empirical distributions of the ratio $m = [v^{3/2}]/[v]^{3/2}$ (left panel), $\eta = V/V_d$ (central panel) and $\xi = N/N_d$ (right panel), all three computed at the daily level for each asset: we randomly group the stocks in equally sized samples and for each of them we compute the empirical distribution respectively of m , η and ξ finding that they are, to a first approximation, stock independent.

Using $p_i \simeq p$, for all the metaorders executed in a day on a stock, the above expression simplifies to:

$$\mathcal{I}_{\text{imp}} = \frac{1}{Y_{\text{imp}} \sqrt{\eta}} \frac{[v]^{3/2}}{[v]^{3/2}} = \frac{1}{Y_{\text{imp}} m \sqrt{\eta}}, \quad (7)$$

where $\eta := V/V_d$ with $V := \sum_{i=1}^N v_i$ is the total ANcerno bet volume, $[\bullet]$ is a daily average operation per stock, and $m > 1$ is the normalized 3/2th moment of v (number of shares of a bet), which depends on the *shape* of the distribution of metaorder size. We have checked that m as well as η are, to a first approximation, independent of the stock (see left and central panels in Fig. 3) indicating that the distribution of metaorder size is, to a large degree, universal and that the ANcerno database is representative of the trading across all stocks. These observations explains why \mathcal{I}_{imp} is also, to a large degree, stock independent.

For the total cost normalisation, our understanding of the invariance property relies on the following empirical fact. The average spread is proportional to the volatility per trade, that is $S = c \sigma_d / \sqrt{N_d}$, where c is a stock independent numerical constant, see [15, 16]. Indeed, the above arguments taken together show that the dimensionless quantity \mathcal{I} can be written as:

$$\mathcal{I} = \frac{1}{Y_{\text{spd}} c \sqrt{\xi} + Y_{\text{imp}} m \sqrt{\eta}}, \quad (8)$$

where $\xi := N/N_d$ is found to be stock independent (see right panel in Fig. 3). Therefore \mathcal{I} is also stock independent. However, the fact that the CV of \mathcal{I} is less than both those of \mathcal{I}_{spd} and \mathcal{I}_{imp} suggests that the Kyle-Obizhaeva “invariant” reflects the fact that metaorders are commensurate to the *total* cost of trading, including both the spread cost and the impact cost.

5 Conclusions

In this work we empirically investigated the *market microstructure invariance hypothesis* recently proposed by Kyle and Obizhaeva [1, 2]. Their conjecture is that the expected dollar cost of executing a *bet* is constant across assets and time. The ANcerno dataset provides a unique laboratory to test this intriguing hypothesis through its available metaorders which can be treated as a proxy for *bets*, i.e. a decision to buy or sell a quantity of institutional size generated by a specific trading idea. Let us summarise what we have achieved in this paper:

- Using *bets* issued for around three thousand stocks, we showed that, at the daily timescale interval, the $N^{3/2}$ scaling law between exchanged risk \mathcal{R} and number of bets is observed independently of the year, the economic sector and the market capitalisation.

- The trading invariant $I := \mathcal{R}/N^{3/2}$ proposed by Kyle and Obizhaeva is non-universal: both its average value $\langle I \rangle$ and its whole distribution clearly depend on the considered stocks, in favour of a *weak universality* interpretation. Furthermore, this quantity has dollar units which makes its hypothesised invariance rather implausible.
- On the basis of dimensional and empirical arguments, we propose a dimensionless invariant defined as a ratio of I and of the bet's total cost, which includes both spread and market impact costs. We find a variance reduction of more than 50%, qualitatively traceable to the proportionality between spread and volatility per trade, and the near invariance of the distributions of bet size, of the volume fraction and number fraction of bets across stocks.

Our empirical analysis has allowed to show that the trading invariance hypothesis holds at the *bet* level in a strong sense provided one considers the exchanged risk *and* the total trading cost of the *bets*. This is in the spirit of Kyle and Obizhaeva's arguments, but takes into account the fact that transaction costs are both asset and epoch dependent. As anticipated in [3], our results strongly suggest that trading "invariance" is a consequence of the endogeneisation of costs in the trading decision of market participants, and has little to do with the Modigliani-Miller theorem. It would actually be quite interesting to investigate other markets such as bond markets, currency markets or futures markets, for which the Modigliani-Miller theorem is totally irrelevant, while trading invariance still holds – at least at the level of single trades [3, 4]. Finally, differences in market structure across countries, such as execution mechanisms, fees and regulations could also challenge the validity of the results presented here.

Acknowledgments

We thank Alexios Beveratos, Laurent Erreca, Antoine Fosset, Charles-Albert Lehalle and Amine Raboun for fruitful discussions. This research was conducted within the *Econophysics & Complex Systems* Research Chair, under the aegis of the Fondation du Risque, the Fondation de l'Ecole polytechnique, the Ecole polytechnique and Capital Fund Management.

Data availability statement

The data were purchased from the company ANcerno Ltd (formerly the Abel Noser Corporation) which is a widely recognised consulting firm that works with institutional investors to monitor their equity trading costs. Its clients include many pension funds and asset managers. The authors do not have permission to redistribute them, even in aggregate form. Requests for this commercial dataset can be addressed directly to the data vendor. See www.ancerno.com for details.

References

- [1] Albert S. Kyle and Anna A. Obizhaeva. Market microstructure invariance: Empirical hypotheses. *Econometrica*, 84(4):1345–1404, 2016.
- [2] Albert S. Kyle and Anna A. Obizhaeva. Dimensional analysis, leverage neutrality, and market microstructure invariance. 2017.
- [3] Michael Benzaquen, Jonathan Donier, and Jean-Philippe Bouchaud. Unravelling the trading invariance hypothesis. *Market Microstructure and Liquidity*, 2(03n04):1650009, 2016.
- [4] Torben G. Andersen, Oleg Bondarenko, Albert S. Kyle, and Anna A. Obizhaeva. Intraday trading invariance in the e-mini s&p 500 futures market. 2016.
- [5] Mathias Pohl, Alexander Ristig, Walter Schachermayer, and Ludovic Tangpi. Theoretical and empirical analysis of trading activity. *arXiv preprint arXiv:1803.04892*, 2018.
- [6] Elia Zarinelli, Michele Treccani, J. Doyne Farmer, and Fabrizio Lillo. Beyond the square root: Evidence for logarithmic dependence of market impact on size and participation rate. *Market Microstructure and Liquidity*, 1(02):1550004, 2015.
- [7] L. Christopher G. Rogers and Stephen E. Satchell. Estimating variance from high, low and closing prices. *The Annals of Applied Probability*, pages 504–512, 1991.
- [8] Bence Tóth, Yves Lempereire, Cyril Deremble, Joachim De Lataillade, Julien Kockelkoren, and Jean-Philippe Bouchaud. Anomalous price impact and the critical nature of liquidity in financial markets. *Physical Review X*, 1(2):021006, 2011.
- [9] Nicolo G. Torre and Mark J. Ferrari. The market impact model. *Horizons, The Barra Newsletter*, 165, 1998.
- [10] Robert Almgren, Chee Thum, Emmanuel Hauptmann, and Hong Li. Direct estimation of equity market impact. *Risk*, 18(7):5862, 2005.
- [11] Robert Engle, Robert Ferstenberg, and Jeffrey Russell. Measuring and modeling execution cost and risk. Chicago GSB Research Paper, no. 08-09, 2006.
- [12] Xavier Brokmann, Emmanuel Serie, Julien Kockelkoren, and Jean-Philippe Bouchaud. Slow decay of impact in equity markets. *Market Microstructure and Liquidity*, 1(02):1550007, 2015.
- [13] Frédéric Bucci, Iacopo Mastromatteo, Zoltán Eisler, Fabrizio Lillo, Jean-Philippe Bouchaud, and Charles-Albert Lehalle. Co-impact: Crowding effects in institutional trading activity. *arXiv preprint arXiv:1804.09565*, 2018.
- [14] Frédéric Bucci, Michael Benzaquen, Fabrizio Lillo, and Jean-Philippe Bouchaud. Crossover from linear to square-root market impact. *arXiv preprint arXiv:1811.05230*, 2018.
- [15] Jean-Philippe Bouchaud, Julius Bonart, Jonathan Donier, and Martin Gould. *Trades, quotes and prices: financial markets under the microscope*. Cambridge University Press, 2018.
- [16] Matthieu Wyart, Jean-Philippe Bouchaud, Julien Kockelkoren, Marc Potters, and Michele Vettorazzo. Relation between bid–ask spread, impact and volatility in order-driven markets. *Quantitative Finance*, 8(1):41–57, 2008.
- [17] Charles M. Jones, Gautam Kaul, and Marc L. Lipson. Transactions, volume, and volatility. *The Review of Financial Studies*, 7(4):631–651, 1994.

A Statistics of metaorder sample

Here we describe some statistics of the metaorders executed from the main investments funds and brokerage firms gathered by ANcerno. The empirical probability distribution of the number of metaorders N per asset, of the risk \mathcal{R}_i exchanged by a metaorder and of the total daily traded risk \mathcal{R} per asset are illustrated in Figure 4. It emerges that both the number of daily metaorders N and the risk measures typically vary over several orders of magnitude. In particular, as evident from the left panel in Fig. 4, there is a significant number of metaorders active every day, since in average ~ 5 metaorders are executed per day for each asset. Furthermore, as shown in the right panel of Fig. 4, both the single bet's risk \mathcal{R}_i and the total daily exchanged risk \mathcal{R} vary over almost eight decades. Note that these statistical properties are approximately independent from the time period and from the economical sector of the asset exchanged through metaorders.




Figure 4: (Left panel) Empirical probability distribution of the daily number N of metaorders per asset: N is broadly distributed over two decades with an average close to 5. (Right panel) Empirical probability distributions of the exchanged risk per metaorder, i.e $\mathcal{R}_i := \sigma_d v_i p_i$, and of the total daily risk per day/assets, i.e $\mathcal{R} := \sum_{i=1}^N \mathcal{R}_i$.

B The 3/2-law under the microscope

One may rightfully wonder whether it is possible to understand the 3/2-law from the statistical properties of the metaorders. To this purpose we start by investigating the individual metaorder's risk \mathcal{R}_i distribution properties as a function of N . We find that when rescaling the metaorder's risk \mathcal{R}_i by the square root of the number N of daily metaorders per asset one obtains a conditional cumulative distribution $\mathcal{P}(\mathcal{R}_i/\sqrt{N}|N)$ dependent on N but with a mean $\mathbb{E}[\mathcal{R}_i/\sqrt{N}|N]$ invariant on N (see Fig. 5). It emerges then that the conditional average metaorder risk \mathcal{R}_i can be predicted from the number N of daily metaorders per asset since $\mathbb{E}[\mathcal{R}_i|N]$ scales as N^γ with $\gamma \simeq 0.5$, that is $\mathbb{E}[\mathcal{R}_i|N] \sim \sqrt{N}$.⁷ It immediately follows that combining this empirical result and the linearity property of the mean, one recovers the 3/2-law $\mathbb{E}[\mathcal{R}|N] \sim N^{3/2}$, since:

$$\mathbb{E}[\mathcal{R}|N] = \mathbb{E}\left[\sum_{i=1}^N \mathcal{R}_i|N\right] = \sum_{i=1}^N \mathbb{E}[\mathcal{R}_i|N] = N \mathbb{E}[\mathcal{R}_i|N] \sim N\sqrt{N} = N^{3/2}. \quad (9)$$

To explain the scaling $\mathbb{E}[\mathcal{R}_i|N] \sim \sqrt{N}$ through the product $\mathbb{E}[\sigma_d|N] \times \mathbb{E}[v_i p_i|N]$ we need to check for the correlation between the daily volatility σ_d and the volume in dollars $v_i p_i$ of a metaorder,

⁷In analogy, the variance $\mathbb{V}[\mathcal{R}_i|N]$ scales linearly with N , i.e. $\mathbb{V}[\mathcal{R}_i|N] \approx \mathbb{E}[\mathcal{R}_i|N]^2$.




Figure 5: Empirical cumulative distribution of the traded metaorder's risk $\mathcal{R}_i = \sigma_d v_i p_i$ without (left panel) and with (right panel) rescaling by the square root of the daily number N of metaorders per asset. The colored vertical lines represent the location of the average for each sample conditional on N . To note that also if the empirical distribution is not an invariant function of N , we observe that $\mathbb{E}[\mathcal{R}_i/\sqrt{N}|N] \simeq \text{const.}$, as evident from the vertical lines in the right panel, which is at the origin of the measured 3/2-law.

which is found to be $\langle \mathbb{C}(\sigma_d, v_i p_i) \rangle \approx 3 \times 10^{-2}$. For each stock we regress $\mathcal{R}_i \sim N^\gamma$, $\sigma_d \sim N^\nu$, $v_i p_i \sim N^\delta$, and we obtain from the empirical distributions of the exponents in Fig. 6 that their average values read $\langle \gamma \rangle = 0.5$, $\langle \nu \rangle = 0.25$ and $\langle \delta \rangle = 0.20$, thus $\langle \gamma \rangle \neq \langle \nu \rangle + \langle \delta \rangle$. However, by looking at the scatter plot of the estimated exponent γ as function of the sum $\nu + \delta$ computed separately for each stock (see bottom right panel in Fig. 6) one observes a clear linear relation.

A possible and intuitive explanation of the non null measured correlation between σ_d and $v_i p_i$ is that metaorders add up to volume, generate market impact and thus increase price volatility. In this way trading volume increases due to both an increase in the number of *bets* and in their sizes, and so does volatility from the increased market impact as discussed for example in [17]. Note that this reasoning is valid even if the metaorders only account for a certain percentage of the total daily market volume $V = \sum_{i=1}^N v_i = \eta V_d$ with η adjusting for the partial view of the ANcerno sample in terms of volume, and for the non-*bet* traded by intermediaries: from our dataset we measure in average $\langle \eta \rangle \approx 5 \times 10^{-2}$.

C Statistics of trading costs

As expected, we find that, for a single *bet* with unsigned volume v , the spread cost $c_{\text{spd}} = S \times v$ is dominant for small volumes, while the market impact cost $c_{\text{imp}} = \sigma_d \times v p \times \sqrt{v/V_d}$ takes over for large volumes (see left panel of Fig. 7). Furthermore, as shown in the right panel of Fig. 7, the average daily market impact cost \mathcal{C}_{imp} accounts on average for $\approx 1/2$ of the total daily trading average cost $\mathcal{C} = \mathcal{C}_{\text{spd}} + \mathcal{C}_{\text{imp}}$, computed using $Y_0 = 3.5$ and $Y = 1.5$ in Eq. (4).




Figure 6: (Top left panel) Empirical distribution of the scaling exponent ν computed for each stock regressing $\sigma_d \sim N^\nu$: in average $\langle \nu \rangle = 0.25$ as shown by the dashed black line. (Top right panel) Empirical distribution of the scaling exponent δ computed for each stock regressing $v_i p_i \sim N^\delta$: in average $\langle \delta \rangle = 0.20$ as shown by the dashed black line. (Bottom left panel) Empirical distribution of the scaling exponent γ computed for each stock regressing $\mathcal{R}_i \sim N^\gamma$: in average $\langle \gamma \rangle = 0.5$ as shown by the dashed black line. (Bottom right panel) Signature scatter plot (coloured by density of data) of the coefficients $\nu + \delta$ and γ respectively estimated conditioning to each stock.




Figure 7: (Left panel) Averaged spread and market impact cost ratios given respectively by c_{spd}/c and c_{imp}/c - with $c_{\text{spd}} = S \times v$ (spread cost), $c_{\text{imp}} = \sigma_d \times v p \times \sqrt{v/V_d}$ (market impact cost) and $c = c_{\text{spd}} + c_{\text{imp}}$ (total cost per *bet*) - as function of the metaorder's order size v/V_d : to note that for a *bet* with small (large) order size the spread (market impact) cost is dominant. (Right panel) Empirical distributions of the $\mathcal{C}_{\text{spd}}/\mathcal{C}$ and $\mathcal{C}_{\text{imp}}/\mathcal{C}$ ratios which give us an idea of the order of magnitude of the different contributions to the total daily average cost per *bet* $\mathcal{C} = \mathcal{C}_{\text{spd}} + \mathcal{C}_{\text{imp}}$ (computed from Eq. 4 fixing $Y_{\text{spd}} = 3.5$ and $Y_{\text{imp}} = 1.5$): the dashed vertical lines represent the location of the mean values equal respectively to $\langle \mathcal{C}_{\text{spd}}/\mathcal{C} \rangle = 0.49$ and $\langle \mathcal{C}_{\text{imp}}/\mathcal{C} \rangle = 0.51$.