Editors' Suggestion | Featured in Physics

Strict Universality of the Square-Root Law in Price Impact across Stocks: A Complete Survey of the Tokyo Stock Exchange

Yuki Sato and Kiyoshi Kanazawa*
Department of Physics, Graduate School of Science, Kyoto University, Kyoto 606-8502, Japan

(Received 13 February 2025; revised 26 June 2025; accepted 22 October 2025; published 15 December 2025)

Universal power laws have been scrutinized in physics and beyond, and a long-standing debate exists in econophysics regarding the strict universality of the nonlinear price impact, commonly referred to as the square-root law (SRL). The SRL posits that the average price impact I follows a power law with respect to transaction volume Q, such that $I(Q) \propto Q^{\delta}$ with $\delta \approx 1/2$. Some researchers argue that the exponent δ should be system specific, without universality. Conversely, others contend that δ should be exactly $1/2$ for all stocks across all countries, implying universality. However, resolving this debate requires high-precision measurements of δ with errors of around 0.1 across hundreds of stocks, which has been extremely challenging due to the scarcity of large microscopic datasets—those that enable tracking the trading behavior of all individual accounts. Here we conclusively support the universality hypothesis of the SRL by a complete survey of all trading accounts for all liquid stocks on the Tokyo Stock Exchange over eight years. Using this comprehensive microscopic dataset, we show that the exponent δ is equal to $1/2$ within statistical errors at both the individual stock level and the individual trader level. Additionally, we rejected two prominent models supporting the nonuniversality hypothesis: the Gabaix-Gopikrishnan-Plerou-Stanley and the Farmer-Gerig-Lillo-Waelbroeck models [Gabaix *et al.*, Nature (London) **423**, 267 (2003).; Q. J. Econ **121**, 461 (2006).; Farmer *et al.*, Quant. Finance **13**, 1743 (2013)]. Our Letter provides exceptionally high-precision evidence for the universality hypothesis in social science and could prove useful in evaluating the price impact by large investors—an important topic even among practitioners.

DOI: 10.1103/65jz-81kv

Introduction—Universality is one of the central topics in physics, whereby common macroscopic laws are observed across various materials [1,2]. For example, universal power laws are observed in critical phenomena, such as $M \propto |T - T_c|^{\delta}$, with an order parameter M near the critical point T_c (e.g., the magnetization M and temperature T for a ferromagnetic system). Interestingly, the exponent δ is independent of the microscopic details of the system. Such a robust exponent is referred to as a *universal exponent* to emphasize the special significance of universality, in contrast to a nonuniversal exponent that may be a fragile observable, as it could easily vary with perturbations to the system's microscopic parameters.

Inspired by such success, physicists have attempted to discover universal power laws beyond physics, including in social science. For instance, in finance, some physicists believe the *square-root law* (SRL) for price impact [3] should be part of such universal scaling laws [4–8]. At the same time, this universality claim has not been fully established for lack of large financial microscopic datasets that enable the tracking of all trading accounts individually. Here we quantitatively support the universality hypothesis through a complete survey of all trading accounts on the Tokyo Stock Exchange (TSE), aiming to conclude the long-standing debate in econophysics.

The SRL is an empirical law that describes the nonlinear price response following large buying (selling) by institutional investors. Consider a trader buying (or selling) Q units of stock by market orders—immediate decisions to buy or sell shares. Following this transaction, the market responds with a price movement Δp. For small Q, a linear response is expected, such that $I(Q) \coloneqq \langle \epsilon \Delta p | Q \rangle \approx \lambda Q$, where $I(Q)$ is called the *price impact* [3], ϵ is the market order sign defined by $\epsilon = +1$ ($\epsilon = -1$) for buying (selling), and $\langle A|B \rangle$ represents the average of A conditioned on B. For large Q, on the contrary, a nonlinear response is empirically observed (see Fig. 1),

$$\mathcal{I}(\mathcal{Q}) \approx c_0 \mathcal{Q}^{\delta} \tag{1}$$

after the nondimensionalization $\mathcal{Q} \coloneqq Q/V_{\mathrm{D}}$ and $\mathcal{I}(\mathcal{Q}) \coloneqq I(\mathcal{Q})/\sigma_{\mathrm{D}}$, where V_{D} is the daily transaction volume and σ_{D} is the daily volatility. Surprisingly, the exponent δ is typically claimed to be close to one-half for any asset worldwide,

*Contact author: kiyoshi@scphys.kyoto-u.ac.jp

 © 2025 American Physical Society

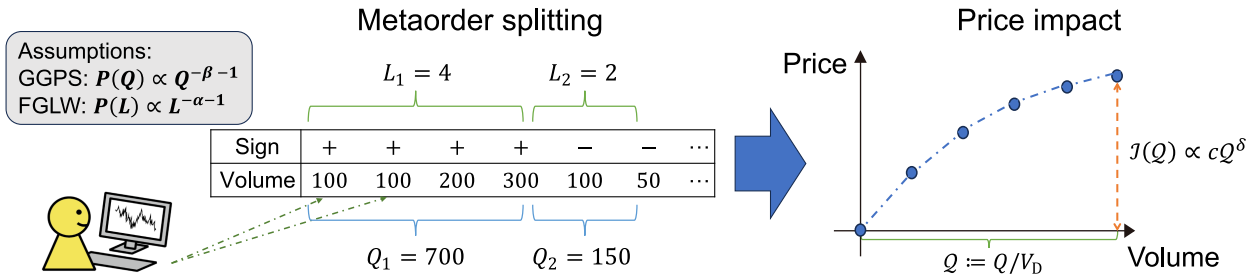


FIG. 1. Square-root price impact law. Practitioners typically split a large volume of market orders (called a metaorder, with size Q) into a long sequence of child orders (with length L). Price impact $I(Q) \coloneqq \langle \epsilon \Delta p | Q \rangle$ is the average mid-price change Δp between the beginning and the end of a buy (sell) metaorder with $\epsilon = +1$ ($\epsilon = -1$) and follows $I(Q) \propto Q^{\delta}$ with $\delta \approx 1/2$.

$$\delta \approx \frac{1}{2}. \tag{2}$$

Thus, the nonlinear scaling (1) is referred to as the SRL [4–8]. For example, the SRL predicts buying 1% of a stock's daily traded shares typically moves the price by about 10% of its daily price range. The SRL closely resembles the universal power law observed in physics and has been an exceptionally appealing topic in econophysics [see Supplemental Material (SM) [9] for a detailed literature review].

However, there has been a long-standing debate regarding the strict universality of the SRL. Some physicists argue that the exponent δ should depend on the microscopic details of financial markets and, therefore, δ should be considered a nonuniversal exponent. For instance, the Gabaix-Gopikrishnan-Plerou-Stanley (GGPS) and the Farmer-Gerig-Lillo-Waelbroeck (FGLW) models are two promising microscopic frameworks that support the nonuniversality hypothesis [24–26]. Conversely, other physicists contend that δ should be exactly one-half for all financial markets in any country, making δ a universal exponent. For example, the latent order-book model [5–8] proposes a minimal mechanism to derive $\delta = 1/2$, supporting the universality hypothesis. Which hypothesis is ultimately correct? This is a major point even in financial economics, as hypothesis selection is strongly relevant for validating and rejecting microscopic models to explain the SRL's mechanisms.

Despite its evident importance, this debate has not yet been settled for several reasons. First, microscopic financial datasets—those that enable tracking trading behavior for all individual trading accounts—are essential to test these hypotheses. Unfortunately, such microscopic datasets are extremely rare, making it challenging to rigorously test these hypotheses. Second, while previous studies have investigated the SRL, their datasets were typically proprietary, provided by specific hedge funds. This means that their datasets are neither comprehensive nor randomly sampled, which introduces the possibility of sampling bias, and the resulting sample sizes are inevitably small. A complete survey—studying all trading accounts for all liquid stocks—has been highly desired to address this issue convincingly. Third, solving these issues requires extremely high accuracy in measuring the exponent δ. To test these hypotheses, we need to measure δ with an accuracy of about 0.1 for hundreds of stocks. Such a high-precision measurement is difficult to achieve in social sciences for lack of large microscopic datasets.

In this Letter, we present conclusive evidence supporting the universality of the SRL by a complete survey of the TSE. We analyzed a large, high-quality microscopic dataset that records the percise trading behavior of all trading accounts on the TSE for all stocks over eight years. Our TSE dataset includes virtual server identifiers (IDs)—a unit of trading accounts on the TSE—enabling us to investigate the resulting price impact Δp associated with the buy (sell) volume Q at the individual account level. Using this dataset, we precisely measured δ by controlling errors of the order of 0.06 and found that δ was equal to $1/2$ for all liquid stocks (more than two thousand data points) within the statistical errors. We also measured δ for all active traders individually and found it to be equal to $1/2$ within statistical errors as well. Furthermore, we rejected two prominent models that had theoretically supported the nonuniversality hypothesis: the GGPS and the FGLW models. Our result supports the universality hypothesis of the SRL with extraordinary accuracy.

Data—Our dataset was provided by Japan Exchange (JPX) Group, Inc., the platform manager of the TSE market. This dataset records the complete life cycle of all order submissions for all stocks on the TSE over almost eight years. The unique feature of our dataset is that it includes virtual server IDs. Virtual server IDs are not technically equivalent to the membership IDs, because any trader may possess several virtual servers. However, we can define effective trader IDs, referred to as *trading desks*, by tracking all the virtual server IDs (see End Matter, Supplemental Material [9], and previous studies [27–29]). In this Letter, we refer to these trading desks as trader IDs for short.

Stock-level histogram of δ—We measured the transaction volume Q and corresponding price impact Δp for all liquid stocks on the TSE. First, we defined trader IDs for each stock and extracted the market-order volume sequences.

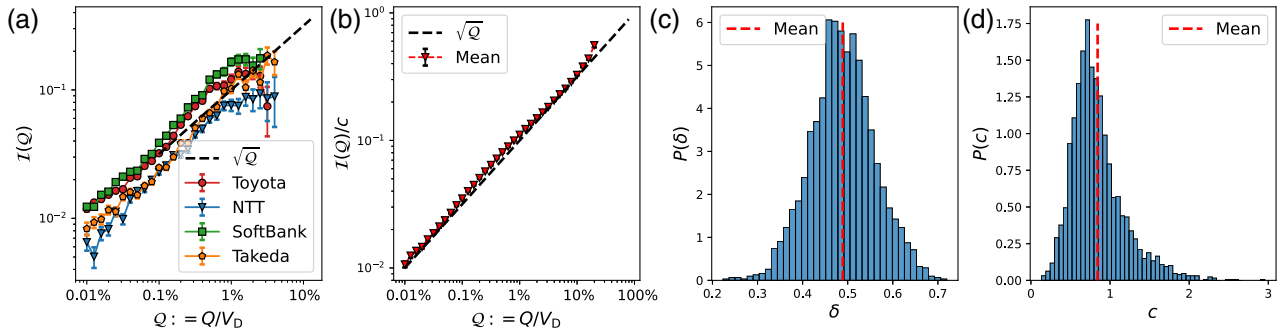


FIG. 2. Stock-level price impact. (a) Average price-impact profile for four stocks [red: Toyota Motor Corp., blue: Nippon Telegraph and Telephone (NTT) Corp., green: SoftBank Group Corp., orange: Takeda Pharmaceutical Company Limited]. Error bars represent the standard errors of the means (SEMs). (b) Aggregate scaling plot, visualizing the average scaling profiles across all stocks, confirms the square-root law for all stocks. (c) Histogram of the exponent δ across all stocks, exhibiting a sharp peak around 0.5 with $\langle\delta\rangle = 0.489$ with its SEM of 0.0015 and $\sigma_\delta := \sqrt{\langle(\delta - \langle\delta\rangle)^2\rangle} = 0.071$. (d) Histogram of the prefactor c across all stocks, with $\langle c\rangle = 0.842$.

Second, we assume that institutional investors typically split their large orders (called metaorders) into small orders (called child orders), as suggested by the literature [28–32]. Specifically, if two successive market orders have the same sign (plus for buy and minus for sell), we assume that both orders are child orders belonging to the same metaorder. The total volume Q is then defined as the sum of all child-order volumes. Third, the price impact Δp is defined as the price movement caused by such a metaorder. In this Letter, liquid stocks are defined as stocks where the total number of metaorders exceeds 10^5.

Figure 2(a) presents the typical price-impact profile $I(Q)$ for four liquid stocks. Additionally, we present Fig. 2(b) as the aggregate scaling plot—visualizing the averaged impact profile across all stocks as one figure. The aggregate scaling plot was created by two steps (see End Matter and [9]): (i) After nondimensionalizing $\mathcal{Q} := Q/V_D$ and $\mathcal{I}(\mathcal{Q}) := I(\mathcal{Q})/\sigma_D$ for each stock, we applied fitting $\mathcal{I}(\mathcal{Q}) = c\mathcal{Q}^{1/2}$ to estimate c. (ii) The average of $\mathcal{I}(\mathcal{Q})/c$ across all stocks was plotted, which collapsed onto a single master curve as theoretically expected. These two figures provide an overall picture of the SRL (1) and qualitatively support its validity across all stocks.

We next present the normalized histogram of δ across all stocks as Fig. 2(c). We estimated δ and c using the power-law scaling (1) by applying the nonlinear relative least squares on binned averages of price impact (see Supplemental Material [9] for a robustness check on the estimation method). Because of the finite-sample size, the δ's error bars for individual stocks are roughly given by $\langle\langle\overline{\sigma_\delta}\rangle\rangle = 0.063$, estimated by a plausible statistical model as discussed in the next paragraph. Figure 2(c) shows a very sharp peak around $1/2$ with the average and the standard deviation,

$$\langle\delta\rangle = 0.489 \pm 0.0015, \quad \sigma_\delta := \sqrt{\langle(\delta - \langle\delta\rangle)^2\rangle} = 0.071, \quad (3)$$

where $\langle A\rangle$ represents the cross-sectional average of A on the TSE dataset across all stocks [33]. This is the first main result of this Letter, suggesting that the exact SRL holds for more than two thousand data points. Also, Fig. 2(d) presents the histogram of the prefactor c. Furthermore, the crossover from the linear to square-root laws was observed consistently with Refs. [4,34] (see Supplemental Material [9]).

Error bar estimation—Statistically, estimating error bars is necessary for δ due to finite sample size. If all the observations of $(\mathcal{Q}, \mathcal{I})$ were independent and identically distributed (IID), simple nonlinear regression would suffice to estimate them. However, our dataset is a time series with serial correlation. For non-IID observations, estimating error bars requires plausible statistical models to account for serial correlation [e.g., the autoregressive integrated moving average (ARIMA) model [35]]. Thus, developing an appropriate statistical model for our dataset is crucial for obtaining reliable error estimates [36].

For this purpose, we develop a simple statistical model with the exact SRL under order splitting and estimate the corresponding error bar (see End Matter and [9]): (i) We randomly shuffled the signs of metaorders, while keeping the physical times of market-order submissions identical to those in our TSE dataset. (ii) The price impacts were assumed to exactly obey the SRL with $\delta = 1/2$, with random contributions. The model parameters were based on our dataset. (iii) We repeatedly generated price time series for 100 iterations and measured δ for all stocks [see Fig. 3(a) for a δ's histogram for one numerical iteration, corresponding to Fig. 2(c) for our TSE dataset]. (iv) The average and standard deviation of δ were evaluated across all Monte Carlo trials as $\langle\langle\overline{\delta}\rangle\rangle$ and $\langle\langle\overline{\sigma_\delta}\rangle\rangle$. For our numerical simulation, $\langle\langle A\rangle\rangle$ represents the Monte Carlo average of A over $N_{MC} = 100$ iterations, whereas $\overline{A}$ represents the average across stocks for one Monte Carlo iteration. This model generates plausible price time series, consistent with diffusive price. Indeed, Ref. [38] develops a nonlinear price-impact model under order splitting, reminiscent of

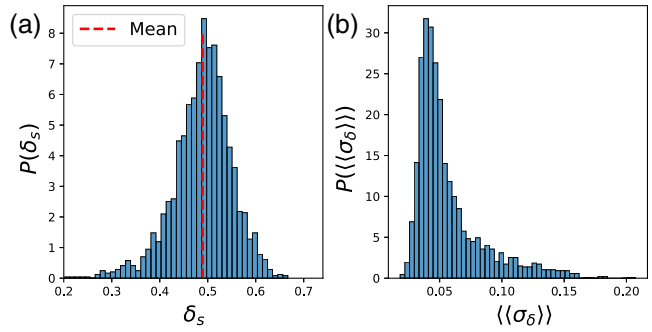


FIG. 3. Error bar estimation based on our statistical model. (a) δ's histogram across all stocks for one Monte Carlo iteration, consistent with the δ's histogram 2(c) for our TSE dataset. (b) Histogram of δ's error bars $\langle\langle\sigma_\delta\rangle\rangle$ for individual stocks by the repeated Monte Carlo simulation. Each error bar was estimated with 100 iterations.

this statistical model, to show its consistency with various stylized facts.

Since $\delta = 1/2$ exactly in this statistical model, the dispersion $\langle\langle\overline{\sigma_\delta}\rangle\rangle$ represents the error bar due to the finite sample size in the frequentist sense. Finally, we obtained

$$\langle\langle\bar{\delta}\rangle\rangle = 0.489 \pm 0.0013, \qquad \langle\langle\overline{\sigma_\delta}\rangle\rangle = 0.063. \tag{4}$$

Hense, we estimate that our δ's error bar is given by $\langle\langle\overline{\sigma_\delta}\rangle\rangle = 0.063$ for the first main result in Fig. 2(c). Also, see Fig. 3(b) for the histogram of error bars for individual stocks.

Remarkably, the cross-sectional average $\langle\delta\rangle$ for the TSE dataset is equal to the Monte Carlo average $\langle\langle\bar{\delta}\rangle\rangle$ for our statistical model within the error bar, such that $|\langle\delta\rangle - \langle\langle\bar{\delta}\rangle\rangle| \ll \langle\langle\overline{\sigma_\delta}\rangle\rangle$. In addition, the cross-sectional dispersion σ_δ for the TSE dataset is almost equal to that $\langle\langle\overline{\sigma_\delta}\rangle\rangle$ for our statistical model, such that $\sigma_\delta \approx \langle\langle\overline{\sigma_\delta}\rangle\rangle$. These facts imply the self-consistency of our statistical analysis. Thus, we concluded that the standard deviation in estimating δ from our dataset is almost attributable to finite sample size, and the universality hypothesis holds for the SRL within statistical errors, at least for the TSE.

Trader-level histogram of δ—While we presented evidence at the stock level, is the SRL universal even at the trader level? In other words, by writing the exponent for the i th trader as $\delta^{(i)}$, is $\delta^{(i)}$ statistically indistinguishable from $1/2$ for all traders, despite the diversity in trading strategies among traders? To answer this question, we studied the price impact of all active traders—roughly defined as traders who submitted metaorders not less than 10^4 times in this Letter (see End Matter and [9]). With this definition, we have 1293 active traders in total. We present the aggregate scaling plot at the individual trader level as Fig. 4(a), qualitatively suggesting the validity of the SRL for all traders. As a second main result, we present the normalized trader-level histogram of $\delta^{(i)}$ as Fig. 4(b). The histogram has clear sharp peak with the average $\langle\delta^{(i)}\rangle$ across all traders on the TSE dataset and standard deviation $\sigma_{\delta^{(i)}} := \sqrt{\langle(\delta^{(i)} - \langle\delta^{(i)}\rangle)^2\rangle}$ as

$$\langle\delta^{(i)}\rangle = 0.493 \pm 0.0050, \qquad \sigma_{\delta^{(i)}} = 0.177, \tag{5}$$

quantitatively supporting the SRL $\langle\delta^{(i)}\rangle \approx 1/2$ within statistical errors. As a reference, we plot the histogram of the prefactor $c^{(i)}$ across all traders in Fig. 4(c).

The standard deviation $\sigma_{\delta^{(i)}}$ in our data was at the same level as that evaluated by our numerical model, $\langle\langle\overline{\delta^{(i)}}\rangle\rangle = 0.521 \pm 0.0048$ and $\langle\langle\overline{\sigma_{\delta^{(i)}}}\rangle\rangle = 0.169$, where $\bar{A}$ represents the average across all traders in our numerical simulation. Since $|\langle\delta^{(i)}\rangle - \langle\langle\overline{\delta^{(i)}}\rangle\rangle| \ll \langle\langle\overline{\sigma_{\delta^{(i)}}}\rangle\rangle$ and $\sigma_{\delta^{(i)}} \approx \langle\langle\overline{\sigma_{\delta^{(i)}}}\rangle\rangle$, we concluded that the universality hypothesis of the SRL is maintained within statistical errors even at the trader level.

Rejecting two nonuniversal models—While we presented direct evidence supporting the universality hypothesis, we now present direct negative evidence against two

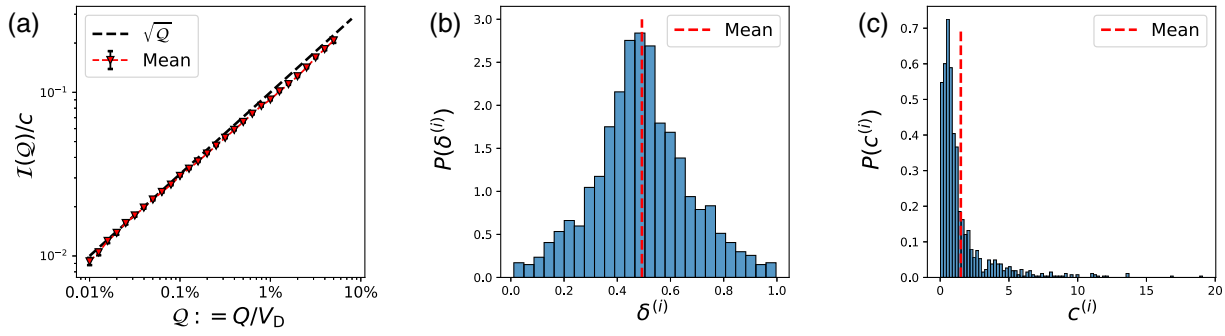


FIG. 4. Trader-level price impact. (a) Trader-level aggregate scaling plot confirms the square-root law for all traders. (b) Histogram of the exponent $\delta^{(i)}$ across all traders for $0 \le \delta^{(i)} \le 1$ shows a sharp peak around 0.5, with $\langle\delta^{(i)}\rangle = 0.493$ with its SEM of 0.0050 and $\sigma_{\delta^{(i)}} := \sqrt{\langle(\delta^{(i)} - \langle\delta^{(i)}\rangle)^2\rangle} = 0.177$. A few outliers were observed outside this region due to fitting errors. (c) Histogram of the prefactor $c^{(i)}$ across all traders, with $\langle c^{(i)}\rangle = 1.501$.

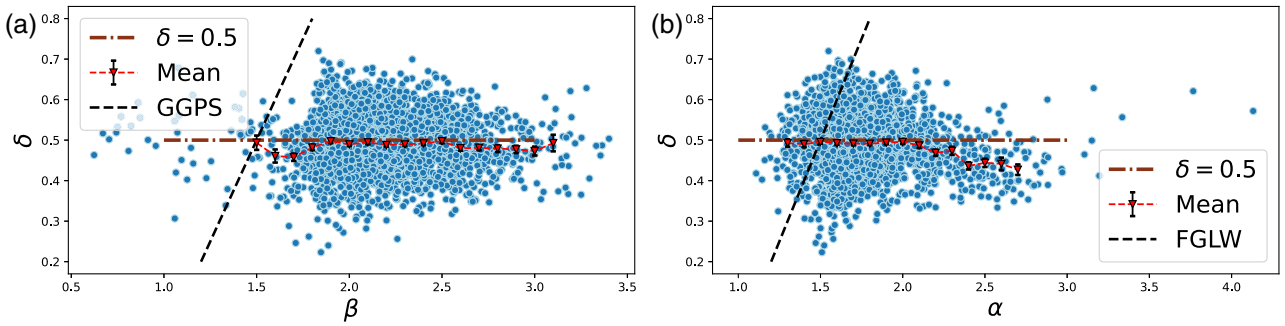


FIG. 5. Quantitative test of (a) the GGPS prediction $\delta = \beta - 1$ and for (b) the FGLW prediction $\delta = \alpha - 1$. No correlation was found and, thus, both models were rejected.

prominent models supporting the nonuniversality hypothesis: the GGPS and the FGLW models [24–26].

The GGPS and FGLW models are based on microeconomic theory in the presence of order splitters: Suppose that an order splitter has a metaorder with total volume Q and splits it into L child orders. Empirically, the distributions of the metaorder volume Q and the number of corresponding child orders L are known to obey power laws [24,25,28–31]: $P(Q) \propto Q^{-\beta-1}$ for large Q and $P(L) \propto L^{-\alpha-1}$ for large L, where β and α are positive power-law exponents characterising individual stocks. Based on this key assumption, the GGPS and FGLW models consider the strategy optimization of market makers in light of inventory risk and fair pricing, respectively. They predict that the exponent δ is given by

$$\text{GGPS}: \delta = \beta - 1, \qquad \text{FGLW}: \delta = \alpha - 1, \tag{6}$$

where Zipf's law is assumed for the company-size distribution as for the GGPS theory.

Interestingly, the typical values of β and α are distributed around 1.5, which is superficially consistent with the SRL [24,25,28,29,31]. However, since β and α can differ from 1.5 depending on the stock, the GGPS and FGLW models have suggested the theoretical nonuniversality of the exponent δ. Their prediction (6) can be tested by precisely estimating β, α, and δ. While a small group of researchers tested this relation (6), both positive and negative results [7,24,25,31] coexist, as direct precise measurement without bias is not an easy task for β, α, and δ, except in a complete survey. We address this issue definitively using our large TSE dataset.

To test the GGPS and FGLW models, it is sufficient to present two high-accuracy scatterplots: one between β and δ, and the other between α and δ. We present Fig. 5 as our final main result. Clearly, there is no correlation at all, rejecting both of the two prominent models that have supported the universality hypothesis, at least on the TSE.

Concluding discussion—This Letter provides the strongest evidence so far supporting the universality hypothesis, at least on the TSE. See also Supplemental Material [9] for thorough statistical robustness of our findings. It would be desirable to perform complete surveys of other markets, which is left as a future study. However, we stress that a single-market survey is sufficient to falsify the universal validity of the prominent nonuniversal models, as shown in this Letter.

Economically, the square-root law implies that market liquidity is larger against metaorder splitting than expected by linear models in traditional economics [39]. Additionally, Ref. [38] implies the SRL plays a key role in validating the efficient market hypothesis—one of the most important concepts in financial economics. Thus, our Letter sheds new light on market liquidity in the presence of large institutional investors.

Our numerical simulation code, used to estimate the finite sample-size effect on the standard deviation of the estimated exponent δ, is available at [40].

Acknowledgments—We appreciate the careful reading of this manuscript by the staff of JPX Group. Also, we thank Jean-Phillipe Bouchaud for his motivating comments on our conference presentations and this manuscript. Y. S. was supported by JSPS KAKENHI (Grant No. 24KJ1328). K. K. was supported by JSPS KAKENHI (Grants No. 21H01560, No. 22H01141, and No. 25K01450) and the JSPS Core-to-Core Program (Grant No. JPJSCCA20200001).

Y. S. developed the program code and conducted the statistical analysis. K. K. designed and managed the project. Both Y. S. and K. K. contributed to writing the manuscript and approved its final content.

We declare no financial conflict of interest. JPX Group provided the original data for this study without financial support.

Data availability—The data that support the findings of this article are not publicly available because they are owned by a third party and the terms of use prevent public distribution. The data are available from the authors upon reasonable request.

[1] H. E. Stanley, *Phase Transitions and Critical Phenomena* (Oxford University Press, New York, 1971).
[2] N. Goldenfeld, *Lecture on Phase Transitions and the Renormalization Group* (CRC Press, Boca Raton, 2019).
[3] J.-P. Bouchaud, J. Bonart, J. Donier, and M. Gould, *Trades, Quotes and Prices: Financial Markets Under the Microscope* (Cambridge University Press, Cambridge, England, 2018).
[4] F. Bucci, M. Benzaquen, F. Lillo, and J.-P. Bouchaud*et al.*, Phys. Rev. Lett. **122**, 108302 (2019).
[5] B. Toth, Y. Lempérière, C. Deremble, J. de Lataillade, J. Kockelkoren, and J.-P. Bouchaud, Phys. Rev. X **1**, 021006 (2011).
[6] I. Mastromatteo, B. Toth, and J.-P. Bouchaud, Phys. Rev. Lett. **113**, 1550008 (2014).
[7] I. Mastromatteo, B. Toth, and J.-P. Bouchaud, Phys. Rev. E **89**, 042805 (2014).
[8] J. Donier, J. Bonart, I. Mastromatteo, and J.-P. Bouchaud, Quant. Finance **15**, 1109 (2015).
[9] See Supplemental Material at http://link.aps.org/supplemental/10.1103/65jz-81kv, which includes Refs. [10–23], for literature review, detailed data description, statistical methods, and robustness check.
[10] R. Almgren *et al.*, Risk **18**, 58 (2005), https://www.cis.upenn.edu/~mkearns/finread/costestim.pdf.
[11] E. Moro, J. Vicente, L. G. Moyano, A. Gerig, J. D. Farmer, G. Vaglica, F. Lillo, and R. N. Mantegna, Phys. Rev. E **80**, 066102 (2009).
[12] F. Bucci, I. Mastromatteo, Z. Eisler, F. Lillo, J.-P. Bouchaud, and C.-A. Lehalle, Quant. Finance **20**, 193 (2020).
[13] F. Bucci *et al.*, Quant. Finance **4**, 193 (2020).
[14] F. Bucci, I. Mastromatteo, M. Benzaquen, and J.-P. Bouchaud, Quant. Finance **19**, 1763 (2019).
[15] E. Bacry *et al.*, Mark. Microstruct. Liq. **1**, 1550009 (2015).
[16] J. Donier and J. Bonart, Mark. Microstruct. Liq. **1**, 1550008 (2015).
[17] X. Brokmann *et al.*, Mark. Microstruct. Liq. **1**, 1550007 (2015).
[18] B. Toth, Z. Eisler, and J.-P. Bouchaud, Wilmott **85**, 70 (2016).
[19] E. Said *et al.*, Mark. Microstruct. Liq. **3**, 1850008 (2017).
[20] E. Said, A. B. H. Ayed, D. Thillou, J.-J. Rabeyrin, and F. Abergel, Quant. Finance **21**, 69 (2021).
[21] J. M. Wooldridge, *Econometric Analysis of Cross Section and Panel Data* (MIT Press, Cambridge, MA, 2018).
[22] F. Capponi and R. Cont, SSRN:3351736 (2019), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3351736.
[23] K. W. Vugrin *et al.*, Water Resour. Res. **43**, W03423 (2007).
[24] X. Gabaix, P. Gopikrishnan, V. Plerou, and H. E. Stanley, Nature (London) **423**, 267 (2003).
[25] X. Gabaix *et al.*, Q. J. Econ **121**, 461 (2006).
[26] J. D. Farmer, A. Gerig, F. Lillo, and H. Waelbroeck, Quant. Finance **13**, 1743 (2013).
[27] K. Goshima, R. Tobe, and J. Uno, Waseda University Institute for Business and Finance Working Paper Series, vol. **19** (2019).
[28] Y. Sato and K. Kanazawa, Phys. Rev. Lett. **131**, 197401 (2023).
[29] Y. Sato and K. Kanazawa, Phys. Rev. Res. **5**, 043131 (2023).
[30] F. Lillo, S. Mike, and J. D. Farmer, Phys. Rev. E **71**, 066122 (2005).
[31] N. Bershova and D. Rakhlin, Quant. Finance **13**, 1759 (2013).
[32] Y. Sato and K. Kanazawa, J. Stat. Phys. **191**, 58 (2024).
[33] If we estimate δ by the nonlinear ordinary least squares on raw data points (without using binned averages), we obtain $\langle \delta_{\mathrm{OLS}} \rangle = 0.474 \pm 0.0012$.
[34] E. Zarinelli, M. Treccani, J. D. Farmer, and F. Lillo, Mark. Microstruct. Liq. **1**, 1550004 (2015).
[35] J. D. Hamilton, *Time Series Analysis* (Princeton University Press, Princeton, NJ, 1994).
[36] If we ignore the non-IID nature of the data and use standard error estimation packages (e.g., `scipy.optimize.least_squares` on Python [37]), the error bar of δ is estimated to be 0.022 via the nonlinear least relative squares on binned averages, or 0.014 via the nonlinear ordinary least squares on raw data points. Both of these values are underestimates of the true error bar in the presence of serial correlation.
[37] P. Virtanen *et al.*, Nat. Methods **17**, 261 (2020).
[38] Y. Sato and K. Kanazawa, arXiv:2502.17906.
[39] A. S. Kyle, Econometrica **53**, 1315 (1985).
[40] https://gitlab.com/Yuki-Sato_JPN/CodeDisclosure-2024-SRL.
[41] A. Clauset, C. R. Shalizi, and M. E. Newman, SIAM Rev. **51**, 661 (2009).
[42] M. Parkinson, J. Bus. **53**, 61 (1980).

End Matter

Appendix A: Data description—We used the JPX order-book dataset on the TSE. These data cover the complete life cycle of all orders, from submissions to transactions or cancellations, for all stocks on the TSE over eight years (from 4th January 2012 to 2nd November 2019).

We divided this dataset into two datasets: one from 4 January 2012 to 18 September 2015 (dataset 1), and another from 24 September 2015 to 2 November 2019 (dataset 2). This is because there was a technical update on the *arrowhead*—the order-matching system managed by the TSE—on 24 September 2015, and all the virtual server IDs were shuffled.

A liquid stock is defined as one where the total number of metaorders exceeds 10^5. This filter is introduced to control the sample size and the resulting accuracy of estimation. In addition, any stock in dataset 1 is regarded as different from stocks in dataset 2, even if their ticker codes are identical. With this definition, the total number of liquid stocks in dataset 1 and dataset 2 are 942 and 1357, respectively. Thus, we have 2299 data points in total regarding liquid stocks.

TABLE I. Schematic of how trader IDs are defined from virtual server IDs.

Virtual server ID	Order ID	Order type	Effective trader ID
V1	O1	Limit order	T1
V2	O1	Cancellation	T1
V3	O2	Limit order	T2
V3	O2	Cancellation	T2
⋮	⋮	⋮	⋮

Our TSE dataset records virtual server IDs corresponding to each order submission. A virtual server acts as an effective unit of trading accounts on the TSE, as all orders are submitted from a virtual server. Technically, a virtual server ID is not exactly technically equivalent to a membership ID, as any membership may possess several virtual servers. In fact, there is a limit on the total number of submissions that can be made during a fixed interval from one virtual server, and high-frequency traders might possess several virtual servers to circumvent this limit. However, we can allocate effective trader IDs by appropriate aggregation of virtual server IDs.

Effective trader IDs are defined as follows: While the lifecycle of a single order ID is typically managed by a single virtual server, there are exceptional cases where multiple virtual servers are involved in the lifecycle of one order. Suppose that virtual server "V1" submits a limit order, and then virtual server "V2" cancels the same limit order under the same order ID "O1" (see Table I). We infer that the same trader operates both virtual servers, V1 and V2, and we assign a unified identifier, "T1," as the effective trader ID. This effective trader ID is referred to as a trader ID for simplicity, while it was originally called trading desks in the literature [27–29].

We define trader IDs for each stock: Even if trader T1 made transactions for two different stocks (such as Toyota and NTT), we define two traders, such as "Dataset1-Toyota-T1" and "Dataset1-NTT-T1."

Appendix B: Definition of metaorders—Metaorders are defined based on the order-splitting hypothesis: typical large investors split a large metaorder into a long sequence of child orders. For example, trader T1 may have a buy (sell) metaorder with a volume of $Q = +100$ ($Q = -100$) and split it into ten buy (sell) child orders, such as $q_1 = q_2 = \ldots = q_{10} = +10$ ($q_1 = q_2 = \ldots = q_{10} = -10$), where the sign $+(-)$ represents a buy (sell) order.

Based on this hypothesis, we assume that any two successive orders belong to the same metaorder if their order signs are identical. For example, if we have a market order sequence $\{q_k^{(i)}\}_k = \{+10, +10, +1, +5, -3, -2, -1, +100\}$ for trader i, the k th metaorder volume $Q_k^{(i)}$ and the corresponding number of child orders $L_k^{(i)}$ are given by $(Q_1^{(i)}, L_1^{(i)}) = (+26, 4)$, $(Q_2^{(i)}, L_2^{(i)}) = (-6, 3)$, and $(Q_3^{(i)}, L_3^{(i)}) = (+100, 1)$. Additionally, any metaorder is defined on a daily basis: market orders on different days are regarded as belonging to separate metaorders.

With this definition, we determined the metaorder volumes and lengths for each stock, and then measured the power-law exponents β and α using Clauset's algorithm [41]. These exponents were used to draw the scatterplots in Fig. 5.

Appendix C: Exceptional handling—The TSE has exceptional rules known as the *daily price limit* and *circuit breaker*, which are aimed at protecting investors from sudden price jumps. Since large price changes are typically caused by significant news, which is beyond the scope of our Letter, we excluded such exceptional days from our data analysis.

This exceptional handling is anticipated to mitigate selection bias in measuring price impact for large Q. For instance, a trader might pause the execution of a buy metaorder if the price goes up unexpectedly fast [3]. Our approach excludes exceptional trading days with sudden price jumps, during which significant implementation bias could exist.

Appendix D: Stock-level price impact—For a given stock, we defined the daily volatility σ_{D} as the difference between the daily maximum and minimum prices [42] (see Supplemental Material [9] for a robustness check using half-day intervals instead of full-day intervals). Also, we defined the daily transaction volume as V_{D}. Using these characteristic daily parameters, we defined the dimensionless volume and corresponding price impact as

$$\mathcal{Q}_k \coloneqq \frac{Q_k}{V_{\mathrm{D}}}, \qquad \mathcal{I}_k \coloneqq \frac{I_k}{\sigma_{\mathrm{D}}} \tag{D1}$$

with the k th metaorder volume Q_k and the corresponding price impact I_k. We then studied the average price impact to apply the nonlinear least-square fitting, such that

$$\mathcal{I}(Q) \coloneqq \langle \mathcal{I}_k | \mathcal{Q}_k = Q \rangle \approx c Q^{\delta} \tag{D2}$$

with the fitting parameters c and δ at the stock level. Our analysis focuses exclusively on metaorder splitting through market orders. Including metaorder splitting via limit orders would introduce estimation bias in the price impact, as the price impact of limit orders is theoretically negative.

The fitting process is given as follows: First, we extracted the metaorders and the resulting price impacts. As a filter, we focused only on metaorders whose liquidation time horizon—defined as the physical time of

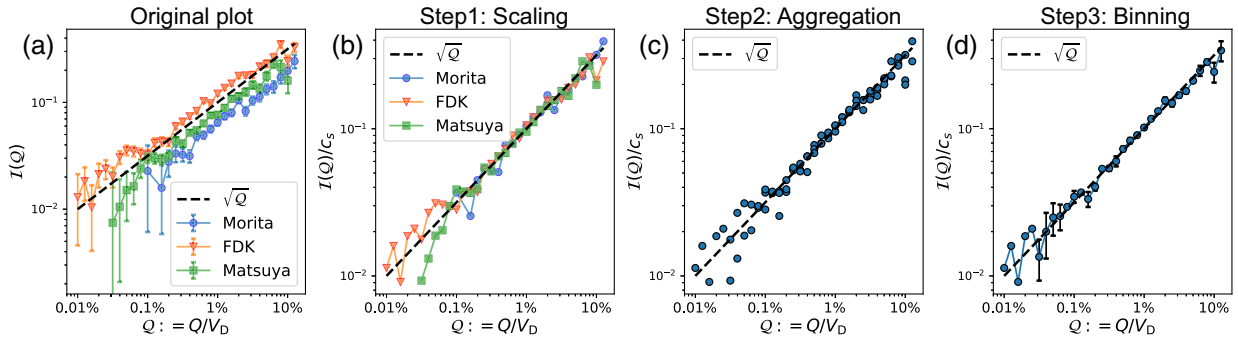


FIG. 6. Schematic example of the aggregate scaling plots: (a) Average price impact profile for the three specific stocks from dataset 2: Morita Holdings Corp. (blue, ticker 6455, $s = 1$), FDK Corp. (orange, ticker 6955, $s = 2$), and Matsuya Co. LTD. (green, ticker 8237, $s = 3$), where s represents the stock index. While these data agree with the square-root law $\mathcal{I}_s(Q) = c_s\sqrt{Q}$ on average, the three coefficient $\{c_s\}_{s=1,2,3}$ are not identical. (b) This heterogeneity is removed by applying the rescaling $\phi(x) = \mathcal{I}_s(Q(x))/c_s$ with $x := \sqrt{Q}$. (c) Next, the average $\phi_{\mathrm{avg}}(x)$ is obtained across all the stocks. (d) Finally, the averaged profile $\phi_{\mathrm{avg}}(x)$ is plotted with bins evenly spaced on a logarithmic scale.

metaorder completion—exceeds 60 sec. This 60-sec rule is inspired by the previous research proposing the 120-sec rule [34]. Second, we introduce logarithmically evenly spaced bins on the Q axis and take the conditional averages for all bins. Third, we focus only on the bins with sufficient sample sizes (more than 100) and apply the fitting to the conditional averages using the method of relative least squares (for the technical details, see Ref. [9]). We thus obtained δ and c for all liquid stocks.

As a robustness check, we also estimated δ and c using ordinary least squares directly on the scatter plot (without relying on binned averages), and the results remained nearly identical (see Ref. [9]). This confirms that our statistical findings are highly robust, irrespective of the measurement method employed.

Appendix E: Trader-level price impact—The method for calculating the trader-level price impact is essentially identical to that used for the stock-level price impact. One important remark is that we focused only on active traders. An active trader is defined as one who meets two conditions: First, the total number of their metaorders must be not less than 10^4. Second, when applying the nonlinear least square fitting to the empirical conditional averages as $\mathcal{I}^{(i)}(Q) \approx c^{(i)}Q^{\delta^{(i)}}$, there must be at least 10 bins where the sample size exceeds 100. This filter was introduced to estimate $\delta^{(i)}$ and $c^{(i)}$ with sufficient statistical accuracy. With this definition, the numbers of active traders were 1293 (415 from dataset 1 and 878 from dataset 2).

Appendix F: Aggregate scaling plot—The aim of the aggregate scaling plot is to visualize a single average scaling master curve across many stocks. While plotting a few scaling curves, as in Fig. 2(a), is the simplest method to qualitatively check the validity of the square-root scaling, it is unrealistic to plot more than 2000 curves. Therefore, we use the aggregate scaling plot as a practical option to show the qualitative validity of the square-root scaling across all stocks.

The generating process of the aggregate scaling plot is as follows (see Fig. 6 for schematic): First, we obtained the price-impact plots $\mathcal{I}(Q)$ for all stocks. We visually checked the validity of the square-root law for all stocks. Second, we applied the fitting $\mathcal{I}(Q) \approx c\sqrt{Q}$ with one parameter c for all stocks and obtained the rescaled profiles $\phi(x) := \mathcal{I}(Q)/c$ with a variable transformation $x := \sqrt{Q}$ across all stocks. Third, while we had many $\phi(x) := \mathcal{I}(Q)/c$ for all stocks, we calculated their average profile as $\phi_{\mathrm{avg}}(x)$. Finally, we plotted $\phi_{\mathrm{avg}}(x)$ at the points that are evenly spaced along the logarithmic horizontal axis.

Appendix G: Numerical simulation—Our nonlinear stochastic model with exact $\delta = 1/2$ is based on our real dataset and random shuffling. (i) All metaorder execution schedules (in physical time) are identical to those in our real dataset. (ii) However, the signs of the metaorders are randomly shuffled using our real dataset. (iii) Furthermore, the price impact of each metaorder is modeled as a Gaussian random variable, following the square-root law on average with the exact half exponent $\delta = 1/2$. The coefficient c and the daily volatility are identical to those in our dataset, but the randomness arises in our model from the shuffling in process (ii) and the Gaussian random variables in process (iii). Since our model has an exact half exponent $\delta = 1/2$, it is feasible to evaluate the estimation error of δ from our real data attributable to the finite sample-size effect, assuming the exact square-root law.