

**Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.**

Microstructure Invariance in U.S. Stock Market Trades

Albert S. Kyle, Anna A. Obizhaeva, and Tugkan Tuzun

2016-034

Please cite this paper as:

Kyle, Albert S., Anna A. Obizhaeva, and Tugkan Tuzun (2016). "Microstructure Invariance in U.S. Stock Market Trades," Finance and Economics Discussion Series 2016-034. Washington: Board of Governors of the Federal Reserve System, <http://dx.doi.org/10.17016/FEDS.2016.034>.

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.

Microstructure Invariance in U.S. Stock Market Trades

ALBERT S. KYLE, ANNA A. OBIZHAEVA, AND TUGKAN TUZUN*

First Draft: September 1, 2010

This Draft: April 19, 2016

This paper studies invariance relationships in tick-by-tick transaction data in the U.S. stock market. Over the 1993–2001 period, the estimated monthly regression coefficients of the log of trade arrival rate on the log of trading activity have an almost constant value of 0.666, strikingly close to the value of $2/3$ predicted by the invariance hypothesis. Over the 2001–14 period, the estimated coefficients rise, and their average value is equal to 0.79, suggesting that the reduction in tick size in 2001 and the subsequent increase in algorithmic trading resulted in a more intense order shredding in more liquid stocks. The distributions of trade sizes, adjusted for differences in trading activity, resemble a log-normal before 2001; there is clearly visible truncation at the round-lot boundary and clustering of trades at even levels. These distributions change dramatically over the 2001–14 period with their means shifting downward. The invariance hypothesis explains about 88 percent of the cross-sectional variation in trade arrival rates and average trade sizes; additional explanatory variables include the invariance-implied measure of effective price volatility.

JEL: G10, G23

Keywords: market microstructure, transactions data, market frictions, trade size, tick size, order shredding, clustering, TAQ data.

* Kyle: Robert H. Smith School of Business, University of Maryland, College Park, MD 20742 USA, akyle@rhsmith.umd.edu. Obizhaeva: New Economic School, Moscow, Russia, aobizhaeva@nes.ru. Tuzun: Board of Governors of the Federal Reserve System, Washington, DC 20551 USA, tugkan.tuzun@frb.gov. We are grateful to Elena Asparouhova, Peter Bossaerts, and Stewart Mayhew for very helpful comments on an earlier draft of this paper. Joseph Saia

Over the past two decades, the U.S. stock market has undergone notable transformations. Technological changes and regulatory reforms have significantly influenced the way stocks are traded. This paper uses market microstructure invariance to define benchmarks for examining how changing market frictions are reflected in cross-sectional and time-series variation in the number and size of trades reported in public transaction data feeds.

Large variation in transaction data makes such analysis challenging. Relying on the benchmarks imposed by the market microstructure invariance of Kyle and Obizhaeva (2016) enables us to filter out the substantial “natural” variation in trading activity across markets and analyze properties of the data due to effects of market frictions.

The invariance hypothesis is based on the intuition that trading in securities markets can be modeled as trading games played at different speeds. Asset managers place bets or meta-orders, which approximately represent uncorrelated decisions to buy or sell specific numbers of shares. A bet may be executed as many smaller orders. The speed with which business time passes is the speed with which new bets are made. In markets for liquid stocks, trading occurs at fast speeds and bets arrive over short horizons, perhaps only a few minutes. In markets for illiquid stocks, trading takes place slowly and bets arrive over longer horizons, perhaps a few months.

The invariance hypothesis conjectures that the dollar risk transferred by bets and the dollar costs of executing bets are the same across markets when measured in *business-time* units corresponding to the rate at which bets occur. This hypothesis implies a specific decomposition of the order flow. *Trading activity*—a measure of aggregate risk transfer—is defined as the product of dollar volume and return volatility. In frictionless markets, invariance implies that the number of bets is proportional to the $2/3$ power of trading activity, and the distribution of bet sizes as a fraction of daily volume is proportional to the negative $2/3$ power of trading activity. These invariance principles define frictionless market benchmarks for examining the number of trades and the distribution of trade sizes in the Trades and Quotes (TAQ) dataset that contains tick-by-tick transactions between 1993 and 2014 for the stocks listed in the U.S. market.

Microstructure invariance is ultimately an empirical hypothesis. Over the 1993–2001 subperiod, a time series of month-by-month regression coefficients of the log of trade arrival rates on the log of trading activity shows that the estimated coefficients remained virtually constant. The estimated coefficient of 0.666 is indeed strikingly close to the benchmark invariance prediction of $2/3$. After 2001, the monthly estimates increase from about 0.690 in 2001 to about 0.77 in 2014; this

provided excellent research assistance. Kyle has worked as a consultant on finance topics for companies, banks, stock exchanges, and various U.S. federal agencies including the Securities and Exchange Commission, the Commodity Futures Trading Commission, and the Department of Justice. He also serves as a non-executive director of a U.S.-based asset manager. The views expressed herein are those of the authors and do not necessarily reflect the views of the Board of Governors or the staff of the Federal Reserve System.

breakdown in the invariance relationships is both statistically and economically significant.

For the years 1993, 2001, and 2014, the empirical distributions of logs of scaled print sizes for stocks sorted into 10 dollar-volume groups and 4 price-volatility groups tell a similar story (see figures 6–10). In 1993, consistent with invariance, all 40 empirical distributions resemble a bell-shaped normal density function with common mean and variance across the 40 subgroups. In 2001 and 2014, the shape of the distributions looks much less like the shape of a normal distribution than in 1993. Furthermore, average scaled trade sizes decrease during the 1993–2014 period by a factor of about 2. Statistical tests clearly reject the hypothesis that scaled trade sizes are distributed as a common log-normal random variable. The rejection arises due to clearly visible microstructure effects such as the one-cent tick size, censoring of trades at the minimum round-lot threshold, and clustering of trades at round-lot sizes such as 100, 1,000, and 5,000 shares, consistent with O’Hara, Yao and Ye (2012) and Alexander and Peterson (2007).

Invariance explains a substantial fraction of the variation in trade arrival rates and average trade sizes across stocks, especially in the first half of the sample. Specifically, when the slope coefficient is restricted to be $\pm 2/3$ as implied by the invariance hypothesis and only intercepts are estimated in separate month-by-month regressions, the time series of R^2 fluctuates around 0.88. Glosten and Harris (1988) find that average trade size (in shares) is negatively related to market depth. Brennan and Subrahmanyam (1998) regress average trade sizes on return volatility, standard deviation of trading volume, market capitalization, number of analysts following a stock, number of institutional investors holding a stock, and the proportion of shares institutional investors hold. The R^2 of 0.92 in their cross-sectional regressions with multiple explanatory variables is only modestly larger than the average R^2 of 0.88 in our restricted regressions. This small difference suggests that other variables offer only limited improvement in explanatory power over the invariance hypothesis.

We attribute the remaining variation to differences in market frictions resulting from how lot size and tick size are related to volume, volatility, and stock price. These market frictions are studied by Harris (1994), Angel (1997), Goldstein and Kavajecz (2000), and Schultz (2000).

A new perspective on these market frictions results from examining the data through the lens of the invariance hypothesis. We introduce two new measures that take into account that trading games run at different speeds in different financial markets. *Effective tick size* is defined as the ratio of tick size to price volatility in business time. *Effective lot size* is defined as a ratio of lot size to a median bet, or equivalently to volume in business time. Both measures are closely related to *effective price volatility*, defined as price volatility in business time. In addition to the 88 percent of the variation in print arrival rates and average print sizes explained by the invariance benchmark, an additional 4.5 percent and 5.5 percent can be attributed to variations in effective price volatility during the 1993–2000

and 2001–14 periods, respectively. It is difficult to further disentangle the effects of lot size and tick size because the effects go in opposite directions.

The invariance hypotheses of Kyle and Obizhaeva (2016) make predictions about bets rather than actual trades or prints resulting from shredding bets into smaller pieces over time and over trading venues via order execution algorithms. Because there is an important distinction between bets and trades, we do not expect estimates based on prints to precisely match the predictions of the invariance hypothesis for bets. Conceptually, the invariance-implied benchmarks are appropriate only under assumptions that the ratio of bet size to trade size and the ratio of intermediate volume to bet volume remain constant across stocks and time. Our results suggest that order shredding and the degree of intermediation may have increased over time, with bets in more liquid stocks recently generating more trades and perhaps more intermediation trades per bet than bets in illiquid stocks.

Our results supplement the findings of Kyle and Obizhaeva (2016), who provide evidence in favor of the invariance hypothesis using a sample of portfolio transitions. Portfolio transition orders are better suited for testing the invariance hypothesis, as they can be thought of as good proxies for bets, but they represent only a small subset of transactions in the U.S. stock market. In contrast, this study is based on a much broader sample that includes all reported transactions in the U.S. stock market; the advantage comes at the expense of having to deal with transaction data affected by order shredding and intermediation.

The remainder of this paper states the implications of the invariance hypothesis, discusses the design of our empirical tests, and presents our results.

1. Testable Implications of the Invariance Hypothesis.

An Invariance Hypothesis. We first review the empirical hypothesis of market microstructure invariance developed in Kyle and Obizhaeva (2016). Market microstructure invariance is based on the intuition that trading in speculative markets can be thought of as the same trading game being played out in different markets at different speeds. This game takes place at a fast speed in active, liquid markets and at a slow speed in inactive, illiquid markets. Asset managers buy and sell securities by placing bets that represent decisions to acquire a long-term position of a specific size, distributed approximately independently from other such decisions. Intermediaries with short-term trading strategies clear markets by taking the other side of bets.

Suppose bets arrive at an expected rate of γ_{jt} bets per day and their size is \tilde{Q}_{jt} shares in asset j and time t . The bet arrival rate γ_{jt} measures the speed of the market. The random variable \tilde{Q}_{jt} has a zero mean; positive values represent buying, and negative values represent selling. Let P_{jt} denote the share price in

dollars, and let V_{jt} denote expected daily volume in shares:

$$(1) \quad V_{jt} = \frac{\zeta_{jt}}{2} \cdot \gamma_{jt} \cdot E|\tilde{Q}_{jt}|.$$

In this equation, expected daily volume is equal to the product of the expected number of bets per day γ_{jt} and their average size $E|\tilde{Q}_{jt}|$, adjusted for the volume multiplier ζ_{jt} .

The parameter ζ_{jt} in equation (1) measures the short-term intermediation trading as the ratio of total volume to bet volume. Non-bet intermediation volume includes trading by market makers, high-frequency traders, and other arbitrageurs who intermediate among long-term bets. The parameter ζ_{jt} intuitively reflects the typical length of intermediation chains in the market; the longer the intermediation chains, the larger the ζ_{jt} . Volume is divided by 2 because each unit of volume has both a buy side and a sell side. If there is no intermediation, then $\zeta_{jt} = 1$. For example, if each bet is intermediated by a single market maker, similar to a New York Stock Exchange (NYSE) specialist intermediating every bet, then $\zeta_{jt} = 2$. If each bet is intermediated by two market makers, who lay off positions trading with one another, similar to Nasdaq dealers in the early 1990s, then $\zeta_{jt} = 3$. If each bet goes through the hands of multiple short-term intermediaries before finding its place in portfolios of long-term traders, then $\zeta_{jt} > 3$.

Let W_{jt} denote trading activity, defined as the product of daily returns volatility σ_{jt} and expected daily dollar volume $P_{jt} \cdot V_{jt}$:

$$(2) \quad W_{jt} = \sigma_{jt} \cdot P_{jt} \cdot V_{jt}.$$

Trading activity measures aggregate risk transfer taking place in the market during the day. It can be easily calculated, as there are usually data available for prices, volume, and volatility. Plugging equation (1) into equation (2) shows that trading activity may be written in terms of less easily observable characteristics of order flow, such as bet arrival rates γ_{jt} and bet sizes $E|\tilde{Q}_{jt}|$:

$$(3) \quad W_{jt} = \frac{\zeta_{jt}}{2} \cdot \sigma_{jt} \cdot P_{jt} \cdot \gamma_{jt} \cdot E|\tilde{Q}_{jt}|.$$

The invariance hypothesis predicts how these characteristics γ_{jt} and \tilde{Q}_{jt} vary across markets with different levels of trading activity W_{jt} .

Business time is measured by the expected arrival rate of bets γ_{jt} , and the dollar risks transferred by bets per unit of business time are the same across assets and time. More specifically, the random variable \tilde{I}_{jt} , defined by

$$(4) \quad \tilde{I}_{jt} := \frac{P_{jt} \cdot \tilde{Q}_{jt} \cdot \sigma_{jt}}{\gamma_{jt}^{1/2}} \stackrel{d}{=} \tilde{I},$$

has an invariant probability distribution \tilde{I} . Here, the risk transferred by a bet per unit of business time is the product of dollar bet size $P_{jt}\tilde{Q}_{jt}$ and returns volatility per unit of business time $\sigma_{jt} \cdot \gamma_{jt}^{-1/2}$.

The invariance hypothesis (4) and equation (3) yield the following testable predictions concerning how bet arrival rate γ_{jt} and bet size \tilde{Q}_{jt} vary with trading activity:

$$(5) \quad \gamma_{jt} = \mu_\gamma \cdot W_{jt}^{2/3},$$

$$(6) \quad \frac{|\tilde{Q}_{jt}|}{V_{jt}} \stackrel{d}{=} \mu_q \cdot W_{jt}^{-2/3} \cdot \tilde{I},$$

where μ_γ and μ_q parameters depend on the volume multiplier ζ_{jt} and the moments of $|\tilde{I}|$.¹ In these equations, the distribution of the random variable \tilde{I} is the same across assets and time. In what follows, we make an identifying assumption that the volume multiplier ζ_{jt} is an invariant constant, and therefore μ_γ does not have indices j and t . Under this assumption, the equations imply that the scaled bet arrival rate $W_{jt}^{-2/3} \cdot \gamma_{jt}$ and distributions of scaled bet sizes $W_{jt}^{2/3} \cdot |\tilde{Q}_{jt}|/V_{jt}$ —that is, all of its percentiles—are invariant across markets. For example, Kyle and Obizhaeva (2016) find that the distribution of logs of scaled bet sizes, $2/3 \cdot \ln(W_{jt}) + \ln(|\tilde{Q}_{jt}|/V_{jt})$, is close to a normal with log-variance 2.53.

Equations (5) and (6) fully describe the composition of the order flow. Changes in trading activity come from both changes in bet sizes and changes in bet arrival rates. Specifically, if ζ and \tilde{I} are invariant and σ_{jt} is constant by assumption, then a 1 percent increase in trading activity W_{jt} is associated with an increase by 2/3 of one percent in the bet arrival rate γ_{jt} and an upward shift by 1/3 of one percent of the entire distribution of unsigned bet sizes $P_{jt}|\tilde{Q}_{jt}|$. The latter is also equivalent to a downward shift by 2/3 of one percent in the distribution of unsigned bet sizes as a fraction of share volume $|\tilde{Q}_{jt}|/V_{jt}$.

Kyle and Obizhaeva (2016) also discuss an invariance hypothesis related to transaction costs. In this paper, we focus only on order flow and leave testing the implications for market impact and bid-ask spreads for future research.

Invariance Implications for TAQ Print Data. The TAQ dataset reports transaction prices and share quantities for all trades in stocks listed in the United States from 1993 to 2014. Each report of a trade execution is called a “print.” We test implications of the invariance hypothesis using data on TAQ print sizes and the number of TAQ prints recorded per day.

Testing invariance this way is not straightforward because prints are different from bets. One bet may generate multiple prints. To minimize transaction costs, traders often break bets or meta-orders into smaller pieces—as documented in Keim and Madhavan (1995), among others—and execute them at several venues, trading with multiple counterparties at multiple prices.

¹The specific parameter values are $\mu_\gamma := E \left[\frac{\zeta_{jt}}{2} \cdot |\tilde{I}| \right]^{-2/3}$ and $\mu_q := E \left[\frac{\zeta_{jt}}{2} \cdot |\tilde{I}| \right]^{-1/3}$.

Let X_{jt} denote the unsigned number of shares in a single print for asset j and at time t . Let ξ_{jt} denote the ratio of the average size of a bet to the average size of a print, so that ξ_{jt} represents the average number of prints per bet. In practice, we expect tiny orders (for example, for the minimum round-lot size of 100 shares) to be executed as one print and gigantic orders to be executed as thousands of small prints or as one print of a gigantic block trade. This multiplier depends on specific details of order-shredding algorithms used by traders—as modeled by Almgren and Chriss (2000) and Obizhaeva and Wang (2013)—and may potentially vary across stocks in a complex and systematic manner that depends on tick size, lot size, and perhaps other factors.

The distribution of average print sizes \tilde{X}_{jt} differs from the distribution of average bet sizes \tilde{Q}_{jt} by a factor ξ_{jt} :

$$(7) \quad \tilde{X}_{jt} = \xi_{jt} \cdot |\tilde{Q}_{jt}|.$$

Let N_{jt} denote the expected number of prints per day for asset j at time t . Each bet \tilde{Q}_{jt} results on average in ξ_{jt} prints and its execution inflates volume by a factor of $\zeta_{jt}/2$ due to induced intermediation volume. The expected number of prints N_{jt} differs from the expected number of bets γ_{jt} by a factor of $\xi_{jt} \cdot \zeta_{jt}/2$:

$$(8) \quad N_{jt} = \xi_{jt} \cdot \frac{\zeta_{jt}}{2} \cdot \gamma_{jt}.$$

It is easy to show that equations (5) and (6) imply the following testable implications for the number of prints and their sizes:

$$(9) \quad N_{jt} = \mu_n \cdot W_{jt}^{\alpha_n},$$

$$(10) \quad \frac{|\tilde{X}_{jt}|}{V_{jt}} \stackrel{d}{=} \mu_x \cdot W_{jt}^{\alpha_x} \cdot \tilde{I},$$

where $\alpha_n = -\alpha_x = 2/3$ and parameters μ_n and μ_x depend on the volume multiplier ζ_{jt} , the order-shredding multiplier ξ_{jt} , and moments of $|\tilde{I}|$.²

As a benchmark for interpreting our empirical results, we make two identifying assumptions. First, assume that there exists an invariant order-shredding multiplier ξ such that $\xi_{jt} = \xi$ for any asset j and time t . Second, assume that there exists an invariant volume multiplier ζ such that $\zeta_{jt} = \zeta$ for any asset j and time t . For simplicity of exposition, results may be interpreted under the identifying assumptions that $\xi = 1$ and $\zeta = 2$. This case corresponds to the hypothesis that each bet is executed as one print against a single intermediary, which makes μ_n

²Specific parameter values are $\mu_n := (\xi_{jt} \cdot \zeta_{jt}/2) \cdot E \left[\frac{\zeta_{jt}}{2} \cdot |\tilde{I}| \right]^{-2/3}$ and $\mu_x := \xi_{jt} \cdot E \left[\frac{\zeta_{jt}}{2} \cdot |\tilde{I}| \right]^{-1/3}$.

and μ_x in equations (9) and (10) constant across markets with no indices j and t . It is also straightforward to write these predictions in the form of regressions.

In actual markets, ξ may certainly deviate from $\xi = 1$ and ζ may also deviate from $\zeta = 2$. The order-shredding algorithms that determine ξ may potentially depend on order sizes and tick size. The amount of intermediation may potentially fluctuate with volatility and trading volume. The effect of relaxing these assumptions is discussed later.

Alternative Hypotheses. We also consider two alternative benchmark hypotheses that make different predictions about the implied exponents on trading activity in equations (9) and (10).

The first alternative hypothesis, *invariance of print frequency*, asserts that the expected number of prints is the same for all stocks, implying $\alpha_n = -\alpha_x = 0$ in equations (9) and (10). Of course, this hypothesis is empirically unrealistic; it is well known that actively traded stocks have more prints than inactive traded stocks. This benchmark is potentially interesting because the illiquidity measure of Amihud (2002) is based on the assumption that order imbalances are proportional to trading volume. This proportionality only holds if the number of bets is constant across markets.

The second alternative hypothesis, *invariance of print sizes*, asserts that the number of prints is proportional to trading activity, implying that $\alpha_n = -\alpha_x = 1$ in equations (9) and (10). Related models are discussed in Tauchen and Pitts (1983), Harris (1987), Jones, Kaul and Lipson (1994), Hasbrouck (1999) and Ané and Geman (2000), among others.

Although extreme, these two hypotheses provide convenient benchmarks for thinking about the relationship between invariance and the existing literature on trade size and trade frequency.

Institutional Details Related to the Microstructure of TAQ Data. Although we make the identifying assumptions that the order-shredding multiplier is invariant ($\xi_{it} = \xi$) and the volume multiplier is invariant ($\zeta_{it} = \zeta$), we do not expect these assumptions to perfectly describe prints in the U.S. equity market. Instead, these assumptions generate benchmarks that can be used to evaluate the economic significance of deviations resulting from changes in various institutional arrangements related to tick size, lot size, and intermediation.

We next discuss how changing institutional arrangements may have affected the order-shredding multiplier and volume multiplier over the 1993–2014 sample period.

Progress in computing technology and regulatory changes have significantly affected trading patterns. As a result, the degree of order shredding varies with order size, ticker symbol, and time.

At the NYSE in the early 1990s, traders typically executed large bets as block trades in the “upstairs” market, in which case at least one side of a reported block trade might correspond precisely to a bet. Prior to the changes in order handling

rules in 1997, Nasdaq dealers often took the other side of entire bets because customers themselves could not place their own orders into a central limit order book and Nasdaq dealers were unhappy if customers “bagged” them by dumping many blocks into the market one after another. As the use of the NYSE’s Direct Order Transfer (DOT) system became more commonly used by professional traders in the 1990s, the use of electronic order submission strategies and order shredding increased. For Nasdaq stocks, this practice accelerated after new order handling rules were implemented in the late 1990s. Bets were shredded more in the second half of our sample (2001–14) than in the second half (1993–2000).

In 2001, the tick size was cut from 6.25 cents (1/16 of a dollar) to one cent. As a result, quoted bid-ask spreads decreased, and fewer shares were shown at the best bid and best offer. Traders used electronic interfaces to place scaled limit orders of small size at adjacent price points separated by one cent, and this led to smaller print sizes for bets of the same size.

Regulation National Market System (NMS), introduced in 2005, further encouraged market fragmentation and competition among multiple trading venues based on speed and efficiency of electronic interfaces, which led to significant order shredding across both time and trading venues. In the past decade, continued improvements in computer technology have widened the use of electronic order handling systems and have made it practical to shred bets for many thousands of shares into tiny pieces of 100 shares or fewer.

To summarize, we expect prints to more closely resemble bets in the earlier part of our sample but to become less representative of bets more recently. We also expect that large bets may result in more prints than small bets.

Several institutional details associated with trade reporting may have further influenced order shredding:

- Although traders may have increasingly shredded orders into “odd lots” of fewer than 100 shares, some traders probably resist shredding orders into numerous odd lots.
- Under NYSE Rule 411(b), broker-dealer member firms have an obligation to consolidate a customer’s odd-lot orders if the share amount of such orders exceeds 100 shares. Other exchanges have similar provisions and have brought enforcement cases against member firms that did not comply with those rules.
- During the sample period, odd-lot transactions were executed through a separate odd-lot trading system, and these small trades were not reported for dissemination on the consolidated tape, as discussed by O’Hara, Yao and Ye (2012).
- “Tape shredding” affects trading patterns. As suggested by Caglio and Mayhew (2012), large orders may be broken up into more trades than small orders to generate additional revenue from sales of consolidated trade and quote data.

- There is some asymmetry in the treatment of buy and sell orders. Large bets are likely to be matched against multiple bets of smaller sizes, also resulting in more TAQ prints for large bets than for smaller ones. According to the Consolidated Tape Association (CTA), the exchanges are required to collect and report last sale data (CTA Plan (1992) Section VII). At the NYSE, for example, it is the duty of the member representing the seller to ensure that a trade has been reported. Because the rules required reporting of “sales” and not “trades,” order splitting may be intrinsically more important for intended buy orders than intended sell orders.

The volume multiplier may also vary with intended order size, ticker symbol, and time. In the beginning of our sample, for example, the volume multiplier was probably larger for orders traded on Nasdaq than for orders traded on the NYSE. Atkyn and Dyl (1997) claim that because Nasdaq dealers were either buyers or sellers in almost every trade at the Nasdaq, the Nasdaq trading volume was inflated by at least a factor of 2 relative to the number of trades actually occurring between end investors. Over time, these patterns may have changed, as dealers’ participation rate in trade facilitation has decreased and trades from other trading systems have begun to be reported on the consolidated tape through the Nasdaq system.

Technological developments have most likely increased the amount of intermediation in securities markets in the second half of the sample. The number of TAQ prints has soared because of order shredding and intermediation by high-frequency traders, who now account for a significant share of volume as described in Chordia, Roll and Subrahmanyam (2011) and Hendershott, Jones and Subrahmanyam (2011). For example, Kirilenko et al. (2015) find that high-frequency traders account for more than 30 percent of stock index futures trading volume but hold their inventories for only a few minutes.

Another important market friction is the tick size. The tick size was reduced from 12.5 cents ($1/8$ of a dollar) to 6.25 cents ($1/16$ of a dollar) in 1997 and to one cent in 2001. Changes in tick size affect trading decisions by changing incentives to provide liquidity and shred orders, as discussed in Harris (1994). When volatility is high and stock price is high, the tick size is small relative to a typical day’s trading range, and thus there are better opportunities for order shredding and making intermediation trades. Although firms can implement stocks splits to adjust percentage tick sizes, these adjustments occur infrequently and with long time lags, as noted in Angel (1997).

As we discuss later, the invariance hypothesis suggests that it is necessary to take into account differences in the speed of trading games across markets to properly analyze the effects of market frictions.

2. Data.

2.1. Data Description

The NYSE TAQ database, which is accessed through Wharton Research Data Services (WRDS), contains trades and quotes reported on the consolidated tape by each participant in the Consolidated Tape Association (CTA) for all stocks listed on all exchanges during the entire 1993–2014 sample period. Because we do not attempt to sign trades as buys or sells based on whether they are executed at the bid or ask price, our analysis employs only data on trades, not quotes. For each trade, the data record the time, exchange, ticker symbol, number of shares traded, execution price, trade condition, and other parameters. The data set contains about 50 billion records with the number of data entries exponentially increasing over time from over 5 million records per month in 1993 to over 400 million records per month in 2014.

We transform the very large raw data files into a smaller dataset convenient for subsequent analysis. First, bad records are removed using standard filters. The TAQ database provides information about the quality of recorded trades using condition and correction codes. We eliminate prints with condition codes 8, 9, *A*, *C*, *D*, *G*, *L*, *N*, *O*, *R*, *X*, *Z* or with correction codes greater than 1. The correction code 8 indicates, for example, that the trade was canceled.

The remaining prints are aggregated in a specific way to reduce the size of the data set while preserving information about the monthly distributions of trade sizes. For each ticker symbol and each day, each print is placed into one of 55 bins constructed based on the number of shares traded. Letting X denote the size of a print in shares, “even” bins correspond to prints of the following exact “even” sizes of $X = 100$, $X = 200$, $X = 300$, $X = 400$, $X = 500$, $X = 1,000$, $X = 2,000$, $X = 3,000$, $X = 4,000$, $X = 5,000$, $X = 10,000$, $X = 15,000$, $X = 20,000$, $X = 25,000$, $X = 30,000$, $X = 40,000$, $X = 50,000$, $X = 60,000$, $X = 70,000$, $X = 75,000$, $X = 80,000$, $X = 90,000$, $X = 100,000$, $X = 200,000$, $X = 300,000$, $X = 400,000$, and $X = 500,000$ shares. “Odd” bins correspond to prints with trade sizes X between adjacent even bins—that is, $X < 100$, $100 < X < 200$, ..., $400,000 < X < 500,000$, and $500,000 < X$. Note that the size of bins grows approximately at a log-rate. Prints with even sizes are considered separately because trades tend to cluster in round-lot sizes.

The result is a much smaller dataset storing the number of trades by day, ticker symbol, and bin. To simplify the subsequent analysis, we make the approximate assumption that the assumed average print size in a bin (in shares) is equal to a midpoint of that bin. If print size is larger than 500,000 shares, we assign it to the 55th bin and assume its size to be 1,000,000 shares. This simplified aggregation makes it possible to capture the most important properties of print-size distributions while implementing our analysis efficiently. The convenience comes, however, at the expense of introducing some additional noise, which may slightly affect results.

For each day and each ticker symbol, the small database also stores the open price; the close price; the number of trades per day; the dollar volume per day; the share volume per day; the close-to-close return; and the volatility, defined as the daily standard deviation of returns over the past 20 trading days from the TAQ data.

For each stock, each bin, and each month, the number of prints is summed over all the days in the month to calculate the frequency of trade sizes. Later we average the frequency distributions of scaled print sizes to construct an empirical distribution of print sizes (in shares) for each stock and each month in the sample. Aggregation by month makes it possible to build better empirical approximations to theoretical distributions for the many inactively-traded stocks with few daily prints.

In addition to calculating the average number of prints per day, we calculate several statistics describing the possibly complicated shape of the distribution of print sizes. We consider the average print size and various percentiles of trade-size distributions. We refer to these percentiles as *trade-weighted percentiles*. For example, the x th trade-weighted percentile corresponds to a print size such that prints with sizes below this threshold constitute x percent of all prints for a given stock in a given month. Note that trade-weighted percentiles effectively put the same weight onto prints of different sizes, which tends to emphasize small trades. For example, if there are 99 prints of 100-share lots and one print of 100,000 shares, then the distribution of print sizes is mostly concentrated at a 100-share level. All trade-weighted percentiles below the 99th percentile are equal to 100 shares. The total trading volume and average print size, however, are largely determined by one big print of 100,000 shares.

Because large trades are economically more important than small trades, we also investigate the right tail of print-size distributions in more detail by examining *volume-weighted percentiles* based on trades' contributions to total volume. The contribution to the total volume by trades from a given print-size bin is calculated based on its midpoint. The volume-weighted distributions give the percentage of trading volume resulting from prints of different sizes. The x th volume-weighted percentile corresponds to a trade size such that trades with sizes below this threshold constitute x percent of total trading volume. In the example in the previous paragraph, percentiles 1–9 are 100 shares and percentiles 10–99 are 100,000 shares.

We report empirical results for both trade- and volume-weighted distributions. Of course, if we know a trade-weighted distribution of print sizes, then we can easily calculate a volume-weighted distribution as well. For the purpose of comparing trade-weighted and volume-weighted distributions, the log-normal is a useful benchmark. It is a straightforward exercise (involving a change in probability measure) to show that if the log of trade-weighted print size is distributed as $N(\mu, \sigma^2)$, then the log of volume-weighted print size is distributed as $N(\mu + \sigma^2, \sigma^2)$. The only difference between the two distributions is the shift in mean; the log-variance remains the same.

To acquire share and exchange codes for stocks in our sample, the monthly data are matched with the Center for Research and Security Prices (CRSP) data, which is accessed through WRDS. Only common stocks listed on the NYSE, American Stock Exchange (AMEX), and Nasdaq from 1993 through 2014 are included in our study. Stocks that had splits or reverse splits in a given month are eliminated from the sample for that month. For each stock and each month, the data are also augmented by adding average daily volume (in dollars and in shares), average price, and the historical volatility. Our final sample includes 1,383,857 stock-month observations. For each of the 263 months between February 1993 and December 2014, there are, on average, observations for about 5,262 stocks.

2.2. Descriptive Statistics

Table 1 describes the data. Panel A reports statistics for the 1993–2000 subperiod. Panel B reports statistics for the 2001–14 subperiod. These statistics are reported separately because the properties of the data changed substantially following decimalization in 2001. Statistics are calculated for all securities in aggregate as well as separately for 10 groups of stocks sorted by average dollar volume. Instead of dividing the securities into 10 deciles with the same number of securities, volume break points are set at the 30th, 50th, 60th, 70th, 75th, 80th, 85th, 90th, and 95th percentiles of dollar volume for the universe of stocks listed on the NYSE with CRSP share codes of 10 and 11. Group 1 contains stocks in the bottom 30th percentile. Group 10 contains stocks in the top 5th percentile. Group 10 approximately corresponds to the universe of S&P 100 stocks. The top five groups approximately cover the universe of S&P 500 stocks. Smaller percentiles for the more active stocks make it possible to focus on the stocks that are the most economically important. For each month, the thresholds are recalculated and stocks are reshuffled across groups.

Summary Statistics before 2001. Panel A of table 1 reports statistical properties of securities and prints in the sample before 2001. For the entire sample of stocks, the average trading volume is \$6.186 million per day, ranging from \$0.15 million for the lowest volume decile to \$176.99 million for the highest volume decile. The average volatility for the entire sample is equal 4 percent per day. The volatility tends to be higher for smaller stocks. The volatility is 4.5 percent for the lowest-volume decile and 2.9 percent for the highest-volume group. Thus, the measure of trading activity, equal to the product of dollar volume and volatility, increases from $0.15 \cdot 0.045$ to $176.99 \cdot 0.029$ —that is, by a factor of 760 from the lowest- to the highest-volume group.

Before 2001, the average print size is equal to \$23,629, ranging from \$11,441 for low-volume stocks to \$89,338 for high-volume stocks. This corresponds to a decrease from 7.6 percent to 0.05 percent of daily volume from the lowest- to the highest-volume group. The median is much lower than the mean, as large prints make the distribution of print sizes positively skewed. The trade-weighted median

ranges from \$5,682 for low-volume stocks to \$28,567 for high-volume stocks, corresponding to a decrease from 3.8 percent to 0.016 percent of daily volume. Note that the invariance hypothesis predicts that the shape of the distributions of trade sizes as a fraction of daily volume must be similar across stocks, with the only difference that their log-means are shifted downward by two-thirds of the increase in a log-trading activity (equation (10)). Because trading activity increases by a factor of 760 from the lowest to the highest deciles, a back-of-the-envelope calculation suggests that the distributions of trade sizes as a fraction of volume should be shifted downward by a factor of $760^{2/3} \approx 80$. This estimate is less than the observed differences in means of $7.6\%/0.05\% \approx 150$ and medians of $3.8\%/0.016\% \approx 240$ between the highest- and lowest-volume groups.

In the 1993–2001 subperiod, the average number of prints recorded per day is 142 for the entire sample, increasing monotonically from 17 to 2,830 from the first to the tenth volume group. The number of prints increases by a factor of $2,830/17 \approx 166$. The invariance hypothesis predicts that the expected number of prints should increase by two-thirds of the increase in trading activity (equation (9))—that is, $760^{2/3} \approx 80$. This back-of-the-envelope calculation suggests that the number of prints increases more than predicted, potentially reflecting a more intensive order shredding in high-volume groups, but further investigation is certainly warranted.

Some print sizes are unusually common in the TAQ data. Before 2001, even-sized trades account for over 61 percent of volume traded and 80 percent of trades executed. The fraction of even-prints is stable across volume groups. The prevalence of these prints validates our choice of bins with even-share bins considered separately. About 16 percent of all transactions and 2 percent of volume traded are executed in 100-share prints. These trades represent 15 percent of transactions for low-volume stocks and 25 percent of transactions for high-volume stocks. There is also a significant number of 1,000-share prints. The large fraction of 1,000-share prints for low-volume stocks relative to high-volume stocks—18 percent versus 14 percent—probably reflects the regulatory rule according to which the Nasdaq market makers had to post quotes for at least 1,000 shares prior to 1997.

Summary Statistics after 2001. Panel B of table 1 describes statistical properties of the sample after 2001. Daily volume more than tripled from \$6.18 million before 2001 to over \$24.4 million after 2001. Volatility decreased from 4 to 3.1 percent per day. These numbers imply that trading activity doubled from the 1993–2001 to 2001–14 subperiods. The average number of prints increased by a factor of 21 from 143 to 3,013, and the average print size decreased by more than a factor of 3 from \$23,629 to \$6,424. Back-of-the-envelope calculations implied by the invariance hypothesis suggest that the changes in print arrival rates and print sizes cannot be explained only by differences in levels of trading activity between the two subperiods but must be attributed to other factors. One of these factors is the order shredding that became increasingly prevalent over time, especially after the reduction in tick size to one cent on January 29, 2001, for NYSE stocks and on April 9, 2001, for Nasdaq stocks. During the 2001–14 subperiod, for example,

100-share trades account for 57 percent of all transactions and 25 percent of volume traded, with these numbers reaching their peaks at 73 percent and 41 percent, respectively, in 2013. Trades of 1,000 shares became less important in this half of the sample. Evident order shredding will significantly affect tests of the invariance hypothesis using TAQ data after 2001.

Frequency and Sizes of TAQ Prints during the 1993–2014 Period. For the period from February 1993 to December 2014, figure 1 plots the 263 monthly values of the scaled mean of the number of prints per month, calculated as $\bar{N}_{m,i} \cdot W_i^{-2/3}$, where $\bar{N}_{m,i}$ is the number of prints per month and W_i is trading activity. Figure 1 also plots the averages of the 20th, 50th, and 80th percentiles of the trade- and volume-weighted distributions of logs of scaled print sizes, calculated as $\ln(|\tilde{X}_i|/V_i \cdot W_i^{2/3})$, where $|\tilde{X}_i|$ is the print size. To facilitate comparison across stocks and across time, all variables are scaled by $W_i^{-2/3}$ and $W_i^{2/3}$ as implied by the invariance hypothesis. The left panel shows the scaled variables averaged across low-volume stocks (group 1). The right panel shows the scaled variables averaged across high-volume stocks (groups 9 and 10).

Trading patterns differ significantly across the 1993–2000 and 2001–14 subperiods. For high-volume stocks, the percentiles of print sizes and print rates do not change much prior to the beginning of decimalization in 2001. Afterward, percentiles of print size decrease steadily, and the average number of prints correspondingly increases. For low-volume stocks, similar changes started to occur even before 2001. Because most low-volume stocks are Nasdaq stocks, the pre-2001 decrease in print sizes and increase in print arrival rates may be explained by the reduction in tick size from 12.5 cents (1/8 of a dollar) to 6.25cents (1/16 of a dollar) announced at Nasdaq in 1997. With the exception of the largest print sizes in high-volume stocks, the downward trend in scaled print sizes and the upward trend in scaled number of prints seems to end at around 2007, and all variables stabilize at some constant levels. A similar pattern can be found in figures of Hendershott, Jones and Subrahmanyam (2011). In the following sections, we will examine these patterns in more detail.

3. Empirical Results

The invariance hypothesis and two alternative models make distinctively different predictions—all three nested in equations (9) and (10)—concerning the differences in the distributions of print sizes and their frequencies across stocks and time. We run our tests based on both the number of prints and distributions of their sizes to determine which of the models provides a more reasonable description of the data.

3.1. Tests Based on Print Frequency

Comparison of Three Models. According to each model, the scaled number of prints $W_i^{-a} \cdot \bar{N}_i$ per day is constant across stocks. The three models differ only

in the exponent a used to normalize the average number of prints. The invariance hypothesis implies $a = 2/3$, the model of invariant print size implies $a = 0$, and the model of invariant print frequency implies $a = 1$.

Figure 2 has three columns and four rows. Each of the three columns contains plots of the log of the average number of prints per day \bar{N} against the log of trading activity W , where the average number of prints is scaled according to each of the three models, respectively. The four rows present results for April 1993 (NYSE and Nasdaq), April 2001, and April 2014. Results are presented for different periods because trading has changed dramatically over time. Also, trading patterns of the NYSE- and Nasdaq-listed stocks differed historically because of differences in regulatory rules across exchanges during the earlier part of the sample, so our results are presented separately for the NYSE- and Nasdaq-listed stocks in April 1993. We choose the month of April to avoid seasonality, as trades tend to cluster much less before the end of the calendar quarter, as shown by Moulton (2005). Each observation corresponds to the average number of prints per day for a given stock in a given month. There are about 6,000 observations on each subplot. If the model is correctly specified, the points are expected to line up along a horizontal line.

In subplots for the invariance hypothesis, observations are scattered around horizontal lines for each of the three years. The invariance hypothesis explains the data very well, especially for the NYSE stocks traded in April 1993. The levels of the horizontal lines move up from the top to bottom figures, showing that the average number of prints has increased over time. For April 1993, the average number of prints is slightly higher for the NYSE stocks than for the Nasdaq stocks. Some Nasdaq stocks with low trading activity have outliers with a very small number of prints. A few of the most illiquid Nasdaq stocks also have outliers with a very high number of prints in April 1993.

In subplots for the model of invariant print frequency, observations are lined up across a line with a positive slope. The model attributes all differences in trading activity entirely to differences in print sizes. Because changes in trading activity are also partially explained by changes in print arrival rates, the model tends to underestimate the number of prints for high-volume stocks and overestimate it for low-volume stocks.

In subplots for the model of invariant print size, observations are lined up along a line with a negative slope. The model attributes all differences in trading activity entirely to differences in trading rates. Because some part of these differences is actually explained by differences in print sizes, the model tends to overestimate the number of prints for high-volume stocks and underestimate it for low-volume stocks.

OLS Estimates of the Number of TAQ Prints. The distinctly different

predictions of the models can be nested into a simple linear regression

$$(11) \quad \ln [\bar{N}_i] = \mu_n + a_n \cdot \ln \left[\frac{W_i}{W_*} \right] + \tilde{\epsilon}.$$

The equation relates the log of the mean number of prints \bar{N}_i per month for stock i to the level of trading activity W_i . The invariance hypothesis predicts $a_n = 2/3$, the model of invariant print frequency predicts $a_n = 0$, and the model of invariant print size predicts $a_n = 1$. For each month, we estimate the parameters μ_n and a_n using an OLS regression, in which there is one observation per stock per month. The constant term μ_n is scaled to represent the log of the expected number of prints for a benchmark stock with trading activity W_* . The scaling constant $W_* = (40)(10^6)(0.02)$ measures trading activity for an arbitrary benchmark stock with a price of \$40 per share, trading volume of 1 million shares per day, and daily volatility of 2 percent per day.

Table 2 presents the results of the regression (11) pooled over time. The six columns show the results for all stocks, the subsets of NYSE-/AMEX-listed stocks, and the subset of Nasdaq-listed stocks, each shown during the 1993–2000 and 2001–14 subperiods. The table reports Fama-MacBeth estimates of the coefficients. To account for trends, Newey-West standard errors are calculated with three lags relative to a linear time trend estimated by OLS regressions from the estimated coefficients $\hat{\mu}_{n,T}$ and $\hat{a}_{n,T}$ for each month. The specific equations from which the linear time trend is estimated are $\hat{\mu}_{n,T} = \mu_{n,0} + \mu_{n,t} \cdot (T - \bar{T})/12 + \tilde{\epsilon}_T$ and $\hat{a}_{n,T} = a_{n,0} + a_{n,t} \cdot (T - \bar{T})/12 + \tilde{\epsilon}_T$, where T is the number of months from the beginning of the subsample and \bar{T} is the mean month in the subsample. For the subperiod February 1993 to December 2000, $T = 1$ for February 1993 and $T = 95$ for December 2000—that is, $\bar{T} = 48$. For the subperiod January 2001 to December 2014, $T = 1$ for January 2001 and $T = 168$ for December 2014—that is, $\bar{T} = 84.5$.

For the 1993–2000 subperiod, the point estimate of $a_{n,0}$ is equal to 0.666, statistically indistinguishable from the predicted value of $2/3$. For the subperiod 2001–2014, the point estimate of $a_{n,0}$ is equal to 0.79. The standard errors of these estimates are 0.002 and 0.005, respectively. Note also that the alternative models predicting $a_{n,0} = 0$ and $a_{n,0} = 1$ are clearly rejected. For the period 1993–2000, the estimated time trend coefficient $a_{N,t}$ of -0.001 per year is statistically insignificant. The estimates $a_{n,0}$ of 0.626 and 0.76 for NYSE stocks during the two sub-periods 1993–2000 and 2001–2014 are smaller than the corresponding estimates $a_{n,0}$ of 0.679 and 0.816 for Nasdaq stocks.

For the 2001–14 period, the estimated time trend coefficient $a_{n,t}$ of 0.007 per year is statistically significant; it approximately corresponds to the increase in a_n by 0.11, from about 0.66 to about 0.77, over the 13-year period.

The point estimate of intercept $\mu_{n,0}$ is equal to 6.147 for the 1993–2000 subperiod and 8.513 for the 2001–14 subperiod, indicating an increase in the average number of prints over time. The constant terms $\mu_{n,t}$ of 0.093 and 0.148 for those subperiods also show statistically significant upward time trends, corresponding to growth

rates in the number of prints per year of about 9.3 percent and 14.8 percent for the two subperiods, respectively.

Figure 3 shows the time series of coefficients from the monthly regressions between February 1993 and December 2014. Panel A presents the time series of 263 month-by-month regression coefficients a_n from monthly regression equations (11). The superimposed horizontal line represents $a_n = 2/3$, the value predicted by the invariance hypothesis. The figure shows that there are two distinctive sub-periods, 1993–2000 and 2001–14. Both the constant term μ_n and the coefficient a_n change over time and may have time trends that are different over the 1993–2000 and 2001–14 subperiods. Over the 1993–2000 subperiod, all estimated coefficients a_n remained virtually constant. The average value of 0.666 is strikingly close to the predicted values of $2/3$. Over the 2001–14 subperiod, the estimates clearly begin to drift up from the values implied by the invariance hypothesis, increasing from about 0.65 to about 0.77 by the end of 2014.

The combined results for both subperiods are consistent with the following interpretation: Over the 1993–2000 subperiod, the invariance hypothesis held because there was a reasonably close correspondence between TAQ prints and bets. Over the 2001–14 subperiod, the invariance relationship broke down because order shredding and intermediation have increased over time and affected high-volume stocks more than low-volume stocks after decimalization in 2001. We discuss these patterns further in section 3.3.

3.2. Tests Based on TAQ Print Sizes

Comparison of Three Models. Next, we examine the trade- and volume-weighted distributions of print sizes scaled for differences in trading activity as suggested by the three models. The three models predict that the distributions of $W^a \cdot |\tilde{X}|/V$ are constant across stocks and time, but the models make different assumptions about the exponent a . The invariance model predicts $a = 2/3$, the model of invariant print frequency predicts $a = 0$, and the model of invariant print size predicts $a = 1$.

To approximate the distribution, we calculate print size $|X|$ based on the midpoint of a print-size bin where it was placed. For each month and for each volume group, the empirical stock-level distributions of scaled print sizes are combined by averaging across stocks in each volume group the frequency distributions of the number of prints in each bin. The results are plotted in figures 4 and 5. For illustrative purposes, only results for April 1993 are presented, but the data for the entire 1993–2014 period are examined more closely below.

Figure 4 shows empirical distributions of logs of scaled print sizes for the NYSE stocks. The figure has three rows and six columns. The three rows contain plots for low-volume stocks in volume group 1, medium-volume stocks in volume groups 2 through 8, and high-volume stocks in volume group 9–10, respectively. The first three columns contain plots of the trade-weighted distributions, with the density of logs of scaled print sizes on the vertical axis. The second three columns contain

plots of the volume-weighted distributions, with the volume contribution of these trades on the vertical axis. In each of the three columns, print sizes are scaled according to the three models. If one of the three models is correct, then the three distributions in the column corresponding to that model should be the same across rows.

To make it easier to interpret results, we superimpose the bell-shaped densities of a normal distribution with the common means and variances equal to the means and variances of trade- and volume-weighted distributions of scaled print sizes based on the entire sample. As previously discussed, if trade-weighted distributions are log-normally distributed, then volume-weighted distributions are log-normally distributed as well. If scaled sizes are distributed as a log-normal, then the three plots in each column of plots are expected to coincide with the superimposed common normal density.

The first column presents the three trade-weighted distributions implied by the invariance hypothesis; they have similar means, variance, and supports. The shapes of empirical distributions bear some resemblance to the superimposed normal density, but the fit is by no means exact. The low-volume group matches the superimposed normal better than the medium- and high-volume groups. The fourth column presents the volume-weighted distributions implied by the invariance hypothesis; they are even more similar to the superimposed common normal density. Thus, the invariance hypothesis explains a substantial part of variation in the distribution of print sizes, especially in the distribution of economically important large trades. The print sizes seem to be distributed similarly to a log-normal; Kyle and Obizhaeva (2016) find that the distribution of portfolio-transition orders is close to a log-normal as well.

For the model of invariant print frequency, the trade-weighted densities are in the second column and the volume-based densities are in the fifth column; they are much less stable across volume groups. In both columns, the distributions shift to the left as trading volume increases, which suggests that the first alternative model understates print sizes for high-volume stocks and overestimates them for low-volume stocks. The model fails to account for the fact that some variation in trading activity is explained by variation in the number of prints.

For the model of invariant print size, the trade-weighted densities are in the third column and the volume-based densities are in the sixth column; they are clearly unstable across volume groups as well. In both columns, the distributions shift to the right as trading volume increases. The second alternative model overstates print sizes for high-volume stocks and underestimates them for low-volume stocks. The alternative models clearly provide worse explanations for the observed variations in print sizes than the invariance hypothesis.

Figure 5 shows our results for the sample of the Nasdaq stocks. Similar to the NYSE stocks, the distributions of print sizes are more stable across volume groups when print sizes are scaled according to the invariance hypothesis. Compared with the NYSE distributions, the Nasdaq distributions are less smooth and have

more spikes, especially the trade-based densities. We attribute these patterns to a regulatory rule that required Nasdaq dealers to quote prices for at least 1,000 shares, leading to a disproportionately large number of 1,000-share Nasdaq trades recorded on the consolidated tape before 1997.

Implications for Log-Normal Distributions. As previously discussed, if the log of trade-weighted scaled print sizes is distributed as $N(\mu, \sigma^2)$, then the log of volume-weighted scaled print sizes is distributed as $N(\mu + \sigma^2, \sigma^2)$. It is interesting to examine how close the log-means and log-variances of the distributions superimposed in figures 4 and 5 for the invariance hypothesis come to satisfying this constraint.

For NYSE stocks, the constraint implies that the volume-weighted mean of 0.97 should be the same as the sum of the trade-weighted mean of -1.15 and its variance of 1.90 . As $-1.15 + 1.90 = 0.75 \neq 0.97$, we see that the constraint fails to hold by a margin of only about 25 percent. The log-variance of 3.04 for the volume-based distribution is much larger than the log-variance of 1.90 for the trade-weighted distribution. This discrepancy is inconsistent with log-normality, which implies that these log-variances should be the same.

For Nasdaq stocks, the volume-weighted mean of 1.2 should be the same as the sum of the trade-weighted mean of -0.19 and its variance of 1.93 . As $-0.19 + 1.93 = 1.72 \neq 1.2$, we see that this constraint fails to hold by a margin of about 30 percent. The log-variance of 1.98 for the volume-based distribution is similar to the log-variance of 1.93 for the trade-weighted distribution, as consistent with the predictions of log-normality. Because these moment restrictions are not perfectly satisfied in the data, the hypothesis of log-normality can be valid only as a very rough approximation at best. Deviations from log-normality include clustering of trades in even-lot sizes (especially prints of 1,000 shares on Nasdaq), censoring and rounding of odd lots, clustering of 100-share trades, and the possibility that very large trades follow a fatter-tailed power-law distribution rather than a log-normal.

OLS Regression Estimates of TAQ Print Sizes, February 1993 to December 2014. We test implications of the invariance hypothesis for print sizes using OLS regressions in which the left-side variable is either a mean or a percentile of either the trade- or volume-weighted distributions of logs of print sizes. For each stock in a given month, the empirical trade- and volume-weighted distributions of logs of print sizes are constructed. Letting $f(\cdot)$ denote a functional that corresponds to either the mean or the p th percentiles (20th, 50th, 80th) of these distributions, these variables are regressed on logs of trading activity:

$$(12) \quad f\left(\ln\left[\frac{\tilde{X}_i}{V_i}\right]\right) = \mu_x + a_x \cdot \ln\left[\frac{W_i}{W_*}\right] + \tilde{\epsilon}_i.$$

There is a clear connection between equations (11) and (12). The expected trading volume is equal to the product of the expected number of prints and expected print size, $V = \bar{N} \cdot E\{\tilde{X}\}$, implying that the left side of equation (11)

is $\log \bar{N} = -\log(E\{\tilde{X}\}/V)$. Thus, the left-hand side variable in equation (12) is similar to reversing the sign on the left-hand-side variable in the regression equation (11). Note that the concavity of the log function implies by Jensen's inequality that the log of the expectation is less than the expectation of the log: $\log(E\{\tilde{X}/V\}) < E\{\log(\tilde{X}/V)\}$. For example, if \tilde{X}_i/V_i were distributed log-normally with the same variance across stocks, then the coefficient estimates for a_γ and a_x would be the same in absolute value but opposite in sign in all of the regressions in equations (12) and (11), but the constant terms μ_γ and μ_x would be different. In our data (as we show below), \tilde{X}_i/V_i deviates from a log-normal distribution sufficiently to make the coefficients a_γ and a_x vary across regression specifications.

Figure 3 shows the time series of coefficients of the monthly regressions from February 1993 to December 2014. There are the time series of 263 month-by-month regression coefficients a_x from regression equation (12) for the 20th, 50th, and 80th percentiles of print sizes over the period 1993–2014. Panel B of figure 3 presents results for trade-weighted distributions. Panel C presents results for volume-weighted distributions. Superimposed horizontal lines represent the level of negative 2/3, the benchmark predicted by the invariance hypothesis.

Again, the figure shows that there are two distinct subperiods. Over the 1993–2000, subperiod all estimated coefficients remained virtually constant. The estimates of a_x are slightly lower than the predicted value of $-2/3$ for the 20th, 50th, and 80th percentiles of trade-weighted distributions; they fluctuate between -0.70 and -0.82 , implying that small print sizes as a fraction of volume decrease faster with trading activity than predicted by the invariance hypothesis. For the 50th and 80th percentiles of volume-weighted distributions, the estimates of a_x fluctuate between -0.45 and -0.70 , somewhat higher than predicted by the invariance hypothesis, and all estimates for the 20th volume-weighted percentiles are very close to negative 2/3.

Over the subperiod 2001–2014, the estimates begin to drift away from their initial levels. For the volume-weighted percentiles, the estimates of a_x decrease from about -0.45 to -0.82 for the 80th percentile, from about -0.54 to -0.89 for the 50th percentile, and from -0.645 to -0.75 for the 20th percentile. For the trade-weighted percentiles, the estimates for the 20th and 50th percentiles do not exhibit any definite patterns, but the estimates for the 80th percentile decrease from about -0.71 to -0.77 . Overall, changes in the large print sizes (right tails) are more significant than changes in the small print sizes (left tails).

Table 3 reports the estimates from regressions in equation (12) pooled over the 1993–2000 period. The first four columns show estimates for the means and percentiles of the trade-weighted distributions. The last four columns show estimates for the means and percentiles of the volume-weighted distributions. Because the monthly estimates of $\hat{\mu}_{x,T}$ and $\hat{a}_{x,T}$ for each month T are changing over time, we add a linear time trend. The table reports Fama-MacBeth estimates of the coefficients, with Newey-West standard errors calculated with three lags relative to a

linear time trend estimated by OLS regressions from the estimated coefficients $\hat{\mu}_{x,T}$ and $\hat{a}_{x,T}$ for each month. As before, the equations estimated for the time trend are $\hat{\mu}_{x,T} = \mu_{x,0} + \mu_{x,t} \cdot (T - \bar{T})/12 + \tilde{\epsilon}_T$ and $\hat{a}_{x,T} = a_{x,0} + a_{x,t} \cdot (T - \bar{T})/12 + \tilde{\epsilon}_T$, where T is the number of months from the beginning of the subsample and \bar{T} is the median month in the subsample.

For the trade-weighted distributions, the estimate of $a_{x,0}$ is equal to -0.741 for the means. This estimate is larger in absolute value by 0.075 than the estimate of 0.666 for the number of prints in table 2. For the trade-weighted percentiles, the estimated coefficients range from -0.781 for the 20th percentile to -0.725 for the 80th percentile. All of these estimates are larger in absolute value than the value of negative $2/3$ predicted by the invariance hypothesis, which implies that print sizes as a fraction of volume tend to decrease with trading activity faster than implied by the theory. For the volume-weighted distributions, the estimated coefficient $a_{x,0}$ is equal to -0.56 for the means; the estimates range from -0.661 for the 20th percentile to -0.481 for the 80th percentile. Across means and percentiles, the standard errors of estimates $a_{x,0}$ range from 0.002 to 0.003 ; these values are similar in magnitude to the averages of standard errors of a_x from the cross-sectional monthly regressions (12), in which those values range from 0.002 to 0.065 . The data suggests that the invariance hypothesis $a_{x,0} = -2/3$ explains the data much better than the alternatives $a_{x,0} = 0$ and $a_{x,0} = -1$, but all three models are statistically rejected.

For the trade-weighted distributions, the estimated time trend coefficient $a_{x,t}$ ranges from 0.008 to 0.012 per year and is statistically significant. For the volume-weighted distributions, the time-trend coefficient is either statistically insignificant or negative.

The estimated intercept $\mu_{x,0}$ of -7.238 in the regression for the trade-weighted means implies that the median print size for the benchmark stock is equal to $\exp(-7.238)$, or 0.07% of daily volume. The estimated intercepts of -8.490 and -6.260 in the regressions for the 20th and 80th percentiles suggest that the average 20th and 80th print-size percentiles are equal to 0.02 percent and 0.19 percent of daily volume for the benchmark stock, respectively. Under the assumption of log-normality, Kyle and Obizhaeva (2016) note that the fraction of volume generated by trades larger than z standard deviations above the log-mean (which equals the median) is given by $1 - N(z - \sigma)$, where σ is the standard deviation for the distribution of the log of trade sizes; based on the trade-weighted variance of 1.90 in figure 4, a log-normality would imply that about 91 percent of volume occurs in print sizes larger than 0.07 percent of daily volume (median trade). The standard errors of $\mu_{x,0}$ in cross-sectional regressions are similar in magnitude and range between 0.014 and 0.031 , across means and percentiles. The negative and statistically significant estimates of time trend $\mu_{x,t}$ indicate that the print sizes have been gradually decreasing during the 1993–2000 subperiod, with a downward drift that is especially pronounced in the right tails of the distributions.

The R^2 is lower in regressions based on volume-weighted distributions than in

regressions based on trade-weighted distributions. For the means, the values of R^2 are 0.90 and 0.69, respectively. The difference in R^2 increases monotonically from a difference between 0.87 and 0.83 for the 20th percentiles to a difference between 0.88 and 0.52 for the 80th percentiles. These numbers show that there is more unexplained variation in large print sizes than in small print sizes. Note that some of this variation may result from the rounding of large odd-size trades to the mid-point of bins or from the small number of observations in the largest bins.

Table 4 reports the estimates for regressions in equation (12) for the 2001–14 subperiod. For the means, the estimates of $a_{x,0}$ are -0.779 for trade-weighted distributions and -0.776 for volume-weighted distributions. These estimates are larger in absolute value than the corresponding estimates of -0.741 and -0.560 for the 1993–2000 subperiod in table 3. All but one of the estimates of -0.757 , -0.769 , and -0.811 for the 20th, 50th, and 80th trade-weighted percentiles, respectively, and -0.796 , -0.843 , and -0.779 for the 20th, 50th, and 80th volume-weighted percentiles, respectively, are also higher in absolute values than the estimates of -0.781 , -0.750 , -0.725 , -0.661 , -0.579 , and -0.481 for the earlier subperiod. The biggest changes occur in the estimates for the 80th percentile of trade-weighted distributions and the 50th and 80th percentiles of volume-weighted distributions, which suggests that recent technological and regulatory changes had the largest effect on the right tail of print-size distributions. The standard errors of $a_{x,0}$ are between 0.003 and 0.004, similar to the averages of standard errors of a_x in monthly regressions (12). This fact validates the adjustment for time trend in the Fama-McBeth procedure; without inclusion of a time trend, the standard errors would have been higher.

The estimates of the intercept $\mu_{x,0}$ for the 1993–2000 subperiod are lower than for the 2001–14 subperiod for the means as well as for all percentiles. For the pooled sample, for example, the estimate of -9.029 in table 4 is lower than the corresponding estimate of -7.238 in table 3; these estimates imply that a typical print size for the benchmark stock fell from 0.07 percent to 0.01 percent of daily volume over the 1993–2000 subperiod. The estimated time trend $\mu_{x,t}$ is negative and statistically significant in all columns, except for the 20th percentile of trade-weighted distributions, also implying that the distributions of print sizes have been shifting downward.

3.3. Detailed Analysis of Market Frictions over Time

We next study how market frictions such as tick size and lot size affect the trading process. More specifically, we examine whether these market frictions can help explain variations across markets in the number and size distribution of prints that cannot be explained by the invariance hypothesis.

To facilitate our analysis, we introduce several new concepts that measure restrictiveness of market frictions in the spirit of the invariance hypothesis. We also discuss why these measures are related to price volatility and share volume in business time as well as to each other.

Effective Price Volatility and Effective Tick Size. Tick size imposes a restriction on the minimum price change. The tick size changed from 1/8 of a dollar to 1/16 of a dollar in the late 1990s and then to one cent in 2001. To assess the restrictiveness of this friction, it is reasonable to compare it with price volatility—that is, with a typical price change. Based on the same intuition, practitioners often measure price volatility in units of tick size. For example, if a \$40 stock has a volatility of 2 percent per day, then daily dollar price volatility is equal to $\$0.80 = \$40 \cdot 0.02$ and a tick size of $\$0.01$ is equal to 1/80 of price volatility.

The invariance hypothesis further suggests that it is more natural to define relative tick size as a fraction of price volatility in business time, not in calendar time. Because business time is proportional to the expected arrival rate of bets γ_{jt} , or, as shown in equation (5), in terms of trading activity it is also proportional to $W_{jt}^{2/3}$, define *effective price volatility* in asset j and time t as

$$(13) \quad \text{Effective Price Volatility}_{jt} := P_{jt} \cdot \sigma_{jt} \cdot \left(\frac{W_{jt}}{W_*} \right)^{-1/3}.$$

This measure is scaled by a constant $W_*^{-1/3}$, previously defined as trading activity in the benchmark stock, so that effective price volatility is exactly equal to daily price volatility for the benchmark stock. In comparison with calendar-time volatility, effective price volatility is lower for more liquid stocks and higher for less liquid stocks to take into account that business time runs faster in more liquid securities.

Now define *effective tick size* as the ratio of the dollar tick size to effective price volatility:

$$(14) \quad \text{Effective Tick Size}_{jt} := \frac{\text{Tick Size}_{jt}}{\frac{P_{jt} \cdot \sigma_{jt}}{P_* \cdot \sigma_*} \cdot \left(\frac{W_{jt}}{W_*} \right)^{-1/3}}.$$

The presence of $P_* \cdot \sigma_*$ in the denominator scales the definition of relative tick size so that it is exactly equal to the dollar tick size for the benchmark stock. A higher effective price volatility makes the effective tick size lower. We conjecture that lower effective tick size may encourage traders to shred meta-orders into a larger number of trades and lead to a larger amount of intermediation.

Effective Share Volume and Effective Lot Size. The lot size imposes a restriction on the minimum number of shares in prints on the tape. For most stocks in the sample, the lot size is equal to 100 shares. An odd lot comprises orders smaller than a round lot (for example, 30 shares) or the non-round-lot portion of larger orders (for example, the 30-share portion of a 530-share order); these odd-lot transactions are executed according to special, often less flexible rules, and information about them is not disseminated to the tape, as described in detail by Hasbrouck, Sofianos and Sosebee (1993).

To assess how restrictive this friction is, one can compare it with the size of a median bet in the corresponding market. Equation (10) implies that median bet

size is also proportional to share volume in business time, which we refer to as *effective share volume*:

$$(15) \quad \text{Effective Share Volume}_{jt} := V_{jt} \cdot \left(\frac{W_{jt}}{W_*} \right)^{-2/3}.$$

This measure takes into account that business time runs faster by a factor of $W_{jt}^{2/3}$ in more liquid securities. It is scaled by the constant $W_*^{-2/3}$ so that effective share volume is exactly equal to daily share volume for the benchmark stock.

Now define *effective lot size* as the ratio of the lot size to effective share volume:

$$(16) \quad \text{Effective Lot Size}_{jt} := \frac{\text{Lot Size}_{jt}}{\frac{V_{jt}}{V_*} \cdot \left(\frac{W_{jt}}{W_*} \right)^{-2/3}}.$$

Practitioners often measure order size as a fraction of daily volume and restrict their trading rate to a fixed fraction—say, 5 percent—of volume in order to control transaction costs. The presence of V_* in the denominator scales the definition of effective lot size so that it is exactly equal to the lot size for the benchmark stock. A lower effective share volume makes the effective lot size larger and therefore more binding, in the sense that a larger fraction of bets falls below that threshold. Some small bets may be executed as odd lots and not recorded on the consolidated tape, some may not be executed at all, and some may be rounded up to round-lot size. The extent of such censoring and rounding is expected to be related to effective lot size.

Because $W_{jt} = V_{jt} \cdot P_{jt} \cdot \sigma_{jt}$, the product of effective tick size (14) and effective lot size (16) is a constant equal to the product of tick size and lot size—that is, the dollar value of one tick on a lot-size trade. It reflects the minimum possible dollar profit (and cost) per transaction in the market. For example, when tick size is one cent and lot size is 100 shares, then this constant is equal to \$1 dollar; if a trader executes at least 100 shares and pays a half spread that is at least one cent, his or her profit (and cost) is at least \$0.50 dollar.

For a given month in the sample, dollar tick size and lot size are usually constant across most of the U.S. stocks (for example, one cent and 100 shares, respectively), implying that effective tick size and effective lot size are closely related to each other and to effective price volatility. In the following analysis, we therefore examine how trading patterns differ across stocks with different levels of effective price volatility.

The effects of the two market frictions on the number of prints and print-size distributions are difficult to separate. For example, higher effective price volatility makes effective tick size lower, but it also makes effective lot size higher; lower effective share volume has similar implications. The first effect, operating through lower effective tick size, probably encourages more intermediation and more order shredding of bets into smaller trades placed at finer adjacent price points as a strategy to avoid front-running, leading to more prints of smaller sizes. The second effect, operating through higher effective lot size, probably encourages more

censoring and rounding up of odd-lot trades, thus leading to fewer prints of larger sizes in the data sample.

Distributions of Print Sizes over Time. We next examine how market frictions affect trading in the U.S. equities market cross-sectionally and through time. After sorting stocks into 10 volume groups and 4 effective price-volatility groups, we analyze distributions of the logs of scaled print sizes $\ln(\frac{X_i}{V_i} \cdot W_i^{2/3})$, with the scaling factor $W_i^{2/3}$ implied by the invariance hypothesis. The invariance hypothesis predicts that these 40 distributions are the same in frictionless markets. To examine how these distributions change over the entire period, we examine the shape of the distributions for the months April 1993, April 2001, and April 2014.

Because dollar tick size and lot size are constant during selected months, sorting stocks into effective price-volatility groups allows us to study the effects of variations in both effective tick size and lot size.

Trade-Weighted Distributions for NYSE-Listed Stocks, April 1993. In April 1993, NYSE-listed stocks were priced in increments of 12.5 cents (1/8 of a dollar). Most of the stocks were traded in multiples of 100-share lots, even though some of them were traded in multiples of 10-share lots, as described in Hasbrouck, Sofianos and Sosebee (1993).

Figure 6 shows the trade-weighted distributions of logs of scaled print sizes for 5 of the 10 volume groups and all 4 effective price-volatility groups for the NYSE-listed stocks in April 1993. The 100-share trades are highlighted in light gray and 1,000-share trades in dark gray. The number of stocks in each subgroup and the average number of trades per day for these stocks are also reported. On each subplot, the density of a normal distribution with the mean of -1.15 and standard deviation of 1.38 is superimposed, calculated for the pooled sample in April 1993. If the invariance hypothesis holds, identification assumptions are valid, and bet sizes are distributed as a log-normal, then all distributions are expected to be invariant across 40 subplots and coincide with the superimposed normal density. For most subgroups, distributions are indeed close to the superimposed normal.

There is, however, a clear truncation below the 100-share odd-lot boundary with clustering of 100-share trades, shown in light gray, in the left tails of the distributions. Because variation in levels of dollar volume within groups is small, the 100-share trades usually fall into the same bin or into two adjacent bins. The only exception is the first volume group, where large variation in trading activity makes the 100-share trades spread over more than four bins.

The empirical distributions have spikes because of clustering of trades at round-lot levels. There are visible spikes at the 100-share level and the 1,000-share level, marked by clustering of light gray and dark gray columns, as well as two spikes in between, corresponding to the 200-share and 500-share levels. For example, there are more trades of 5,000 shares than 4,000 or 6,000 shares, and far more than 4,900 or 5,100 shares. For the subsample of stocks with low volume and low price

volatility, large variation in trading activity smooths out spikes in the distribution of print sizes.

A visual inspection suggests that, holding effective price volatility fixed, the supports of distributions stay relatively constant across volume groups, but their shapes become more skewed to the right as volume increases, especially when price volatility is low, consistent with large orders being shredded into smaller trades.

Holding dollar volume fixed, the distributions vary across effective price-volatility groups in a systematic way as well. When effective price volatility increases, the 100-share boundary becomes more binding, the truncation threshold shifts to the right, and the effects of censoring and rounding to the 100-share boundary become more pronounced. At the same time, the relative tick size decreases, thus encouraging more order shredding. The first effect seems to dominate, because the average number of prints decreases with price volatility. For high-volume stocks, for example, the average number of prints decreases monotonically from 1,139 prints recorded per day for low-volatility stocks to only 285 prints for high-volatility stocks. In the absence of any market frictions, the number of trades is expected to be relatively constant within a given volume group because volatility does not vary much.

Volume-Weighted Distributions for NYSE-Listed Stocks, April 1993.

Figure 7 shows the volume-weighted distributions of logs of scaled print sizes for the NYSE stocks in April 1993. In comparison with the trade-weighted distributions in figure 6, the volume-weighted distributions put more weight onto larger trades and allow us to see more clearly the distribution of large print sizes.

According to Hasbrouck, Sofianos and Sosebee (1993), a block of 10,000 shares was not a particularly large trade in 1993. Some block trades (of 10,000 shares or more), especially block trades in illiquid stocks, were executed in the upstairs market. The average size of upstairs-facilitated blocks was about 43,000 shares.

In comparison with the trade-weighted distributions in figure 6, the volume-weighted distributions are more stable across subgroups and more closely resemble the superimposed normal distribution. On most plots, the space below the bell-shaped density function is filled up. The truncation at the odd-lot boundary is almost invisible, because the numerous 100-share trades almost “disappear” from the left tail of the distribution, as they contribute little to overall volume.

Small gaps in the distributions relative to a log-normal can be seen in mid-range print sizes between 1,000 shares and 10,000 shares. Perhaps these gaps represent intended orders shredded into smaller trades. The strong visual resemblance of the graphs to a log-normal, as in Kyle and Obizhaeva (2016), are consistent with the interpretation that most of the largest orders in 1993 appear to have been executed as single blocks, generating large print sizes. An exception is the low-volume group for which the largest orders appear to be shredded because the distributions are skewed to the right. Although the volume-weighted distributions are much smoother than the trade-weighted distributions, small spikes are still detectable. These spikes, which likely correspond to clusters of trades at the levels

of 1,000 shares, 5,000 shares, and 10,000 shares, are clearly visible, for example, in distributions for stocks with high volume and high effective price volatility. There are also a few spikes in the far-right tails of several distributions, suggesting that a few very large prints occur in the data more often than explained by log-normality.

Trade-Weighted Distributions for Nasdaq-Listed Stocks, April 1993.

In April 1993, the Nasdaq-listed stocks usually had the minimum lot size of 100 shares. Quotes were restricted to increments of $1/8$ of a dollar if the bid price exceeded \$10.00, but trades were permitted in finer increments of $1/64$ of a dollar for all stocks, even though these prices were then rounded to the nearest eighths for reporting, as described in Christie, Harris and Schultz (1994) and Smith, Selway and McCormick (1998).

Also, in 1988, the Securities and Exchange Commission required Nasdaq market makers to have a quotation size of at least 1,000 shares for most stocks. The rule mostly affected large stocks, and, indeed, we observe larger spikes in subplots for high-volume stocks. For small stocks, the rule was slightly different. For example, orders smaller than 1,000 shares could be executed through the Small Order Execution System (SOES) in stocks that were trading at prices lower than \$250 per share. After 1996, the minimum quote size restriction was gradually removed. Under the Actual Size Rule, the minimum quote size was reduced from 1,000 to 100 shares, first for 50 pilot stocks in January 1997, then for an additional 104 stocks in November 1997, and finally for all others.

Figure 8 shows the trade-weighted distributions of logs of scaled print sizes for the Nasdaq stocks in April 1993. The biggest difference between the trade-weighted distributions of the Nasdaq stocks and the NYSE stocks is the very large fraction of 1,000-share trades, shown as dark gray spikes, typically in the middle of the Nasdaq distributions. These spikes can be attributed to the requirement to quote at least 1,000 shares. In line with this explanation, we do not observe the clustering at the 1,000-share level after 2001 (unreported). Apart from the clustering in the 1,000-share level and truncation at the 100-share level, the distributions bear some resemblance to the superimposed normal distribution.

Trade-Weighted Distributions for All Stocks, April 2001 and 2014.

After decimalization in 2001, U.S. stocks traded in increments of one cent. The round-lot size is 100 shares; odd-lot trades account for a large fraction of trades, but they are not reported in the TAQ dataset. The U.S. stock market is fragmented into multiple trading venues. Angel, Harris and Spatt (2011, 2015) discuss other recent innovations in trading.

Figures 9 and 10 show the trade-weighted distributions of logs of scaled print sizes for all stocks traded in April 2001 and April 2014, respectively. In 2001, decimalization and use of electronic interfaces led to a significant increase in order shredding, the effect of which is clearly seen in both figures.

The frequency of trades has increased significantly over time. For high-volume and low-volatility stocks, for example, there were, on average, only 325 trades

per day in April 1993, increasing to 16,230 trades per day in April 2001 and 41,295 trades per day in April 2014. The distributions of scaled print sizes have shifted substantially to the left during the 1993–2014 period. Based on the means of superimposed normals, for example, the median print size dropped from 0.11 percent of daily volume for the NYSE stocks in April 1993 and 0.08 percent for the Nasdaq stocks to 0.003 percent in 2001 and only 0.001 percent in 2014. The market for block trades seems almost to have disappeared, and trading is now dominated by transactions of 100 shares. Indeed, trades of 100 shares constitute about 57 percent of trades executed and 38 percent of volume traded in 2014.

Note that the estimated log-variances of 2.05 in 1993, 1.78 in 2001, and 1.21 in 2014 for trade-weighted distributions of scaled print sizes are lower than the variance of 2.50 for the distributions of portfolio-transition orders, reported by Kyle and Obizhaeva (2016). A gradual decrease in log-variance is consistent with the hypothesis that large bets in liquid stocks in 2014 result in disproportionately more prints than small bets in illiquid stocks.

Regressions with Effective Price Volatility. Table 5 presents Fama-MacBeth estimates of μ_γ and a_σ from monthly regressions

$$(17) \quad \ln [\bar{N}_i] = \mu_n + \frac{2}{3} \cdot \ln \left[\frac{W_i}{W_*} \right] + a_\sigma \cdot \ln \left[\frac{P_i \cdot \sigma_i}{P_* \cdot \sigma_*} \cdot \left(\frac{W_i}{W_*} \right)^{-1/3} \right] + \tilde{\epsilon}_i.$$

The regression effectively imposes the invariance restriction of $a_n = 2/3$ in regression (11) and adds effective price volatility as an additional explanatory variable to capture the effect of both market frictions. The table reports Fama-MacBeth estimates of the coefficients, with Newey-West standard errors calculated with three lags relative to a linear time trend estimated by OLS regressions from the estimated monthly coefficients $\hat{\mu}_{n,T}$ and $\hat{a}_{\sigma,T}$ for each month. Specifically, the specification is $\hat{\mu}_{n,T} = \mu_{n,0} + \mu_{n,t} \cdot (T - \bar{T})/12 + \tilde{\epsilon}_T$ and $\hat{a}_{\sigma,T} = a_{\sigma,0} + a_{\sigma,t} \cdot (T - \bar{T})/12 + \tilde{\epsilon}_T$, where T is the number of months from the beginning of the subsample, and \bar{T} is the mean month in the subsample. The six columns show the results for the entire sample as well as subsets of NYSE/AMEX-listed stocks and Nasdaq-listed stocks during the 1993–2000 and 2001–14 subperiods.

The point estimates for $a_{\sigma,0}$ are negative and statistically significant for all subsamples. The estimates of -0.471 , -0.338 , and -0.497 for the 1993–2000 subperiod are smaller in absolute terms than the corresponding estimates of -0.608 , -0.475 , and -0.676 for the 2001–14 subperiod. The standard errors range from 0.003 to 0.001. The point estimates of $a_{\sigma,t}$ are equal to -0.007 , -0.02 , -0.007 , -0.027 , -0.023 , and -0.013 , with standard errors between 0.001 and 0.003. There is an inverse relationship between effective price volatility and the number of prints. Higher effective volatility implies fewer prints in the context of the invariance hypothesis. The estimates of $\mu_{n,0}$ and $\mu_{n,t}$ are not too different from the corresponding estimates in table 2.

The significant increase in values of R^2 in regressions (17) relative to values of R^2 in regression (11) constrained with $a_n = 2/3$ shows that the cross-sectional variation in the number of prints, unexplained by the invariance hypothesis, can be partially attributed to differences in effective price volatility. For the entire sample, adding effective price volatility as an explanatory variable increases the R^2 from 0.873 to 0.918 for the 1993–2000 subperiod and from 0.899 to 0.955 for the 2001–14 subperiod. For NYSE stocks, the R^2 increases from 0.908 and 0.917 to 0.924 and 0.955; for Nasdaq stocks, the R^2 increases from 0.857 and 0.881 to 0.917 and 0.959 during the two subperiods.

Finally, we analyze the R^2 in the regression that imposes the invariance restriction $a_n = 2/3$ in regression (11), but allows the coefficients on the three components of trading activity W_i (volume V_i , price P_i , and volatility σ_i) to vary freely:

$$(18) \quad \ln [\bar{N}_i] = \mu_n + \frac{2}{3} \ln \left[\frac{W_i}{W_*} \right] + b_1 \cdot \ln \left[\frac{V_i}{(10^6)} \right] + b_2 \cdot \ln \left[\frac{P_i}{(40)} \right] + b_3 \cdot \ln \left[\frac{\sigma_i}{(0.02)} \right] + \tilde{\epsilon}.$$

For the entire 1993–2014 period, the estimates are $\hat{b}_1 = 0.20$ for the coefficient on volume V_i , $\hat{b}_2 = -0.31$ for the coefficient on price P_i , and $\hat{b}_3 = -0.5$ for the coefficient on volatility σ_i (not reported). All coefficients are statistically different from zero. Similar patterns are observed for other subperiods and subsamples of the NYSE stocks and the Nasdaq stocks. The exponents for volatility behave similarly to the exponents for price and differently from the exponents for volume, thus implying that the rejection of the invariance hypothesis might depend in a subtle manner on how effective price-volatility influences incentives to shred orders and make intermediation trades. Note that for the pooled sample, the R^2 increases from 0.873 to 0.928 for the 1993–2000 subperiod and from 0.899 to 0.970 for the 2001–14 subperiod. The R^2 of 0.928 and 0.970 are only slightly larger than the R^2 of 0.918 and 0.955 in regressions (17), respectively. Although statistically significant, the addition of two extra degrees of freedom beyond the effective price volatility improves the R^2 by only a small amount.

4. Conclusion

The distributions of TAQ print sizes (adjusted for trading activity as suggested by invariance hypotheses) resemble a log-normal, with truncation below the 100-share odd-lot boundary. The resemblance was stronger during the earlier 1993–2001 period than during the later 2001–14 period, and it shows up more clearly in volume-weighted distributions than in trade-weighted distributions.

The invariance hypothesis explains about 88 percent of variation across stocks in the number of prints. The unexplained 10 percent can be most likely attributed to other microstructure effects such as order shredding, intermediation activity, and various market frictions like lot size and tick size. For example, subtle effects of effective price volatility explain an additional 4.5 percent and 5.5 percent of variations during the 1993–2000 and 2001–14 periods, respectively. An interesting

topic for future research would be to analyze these effects at a deeper level by designing more refined econometric tests.

REFERENCES

- Alexander, Gordon, and Mark Peterson.** 2007. “An Analysis of Trade-Size Clustering and its Relation to Stealth Trading.” *Journal of Financial Economics*, 84: 435–471.
- Almgren, Robert, and Neil Chriss.** 2000. “Optimal Execution of Portfolio Transactions.” *Journal of Risk*, 3: 5–39.
- Amihud, Yakov.** 2002. “Illiquidity and Stock Returns: Cross-Section and Time-Series Effects.” *Journal of Financial Markets*, 5(1): 31–56.
- Ané, Thierry, and Helyette Geman.** 2000. “Order Flow, Transaction Clock, and Normality of Asset Returns.” *Journal of Finance*, 55(5): 2259–2284.
- Angel, James J.** 1997. “Tick Size, Share Prices, and Stock Splits.” *Journal of Finance*, 52(2): 655–681.
- Angel, James J., Lawrence E. Harris, and Chester S. Spatt.** 2011. “Equity Trading in the 21st Century.” *The Quarterly Journal of Finance*, 1(01): 1–53.
- Angel, James J., Lawrence E. Harris, and Chester S. Spatt.** 2015. “Equity Trading in the 21st Century: An Update.” *The Quarterly Journal of Finance*, 5(01): 1550002.
- Atkyn, Allen B., and Edward A. Dyl.** 1997. “Transactions Costs and Holding Periods for Common Stocks.” *Journal of Finance*, 52(1): 309325.
- Brennan, Michael, and Avanidhar Subrahmanyam.** 1998. “The Determinants of Average Trade Size.” *Journal of Business*, 71(1): 1–25.
- Caglio, Cecilia, and Stewart Mayhew.** 2012. “Equity Trading and the Allocation of Market Data Revenue.” FEDS Working Paper No. 2012-65.
- Chordia, Tarun, Richard Roll, and Avanidhar Subrahmanyam.** 2011. “Recent Trends in Trading Activity and Market Quality.” *Journal of Financial Economics*, 101: 243–263.
- Christie, William G., Jeffrey H. Harris, and Paul H. Schultz.** 1994. “Why Did NASDAQ Market Makers Stop Avoiding Odd-Eighth Quotes?” *Journal of Finance*, 49(5): 1841–1860.
- Glosten, Lawrence, and Lawrence Harris.** 1988. “Estimating the Components of the Bid-Ask Spread.” *Journal of Financial Economics*, 21: 123–142.

- Goldstein, Michael A., and Kenneth A. Kavajecz.** 2000. "Eighths, Sixteenths, and Market Depth: Changes in Tick Size and Liquidity Provision on the NYSE." *Journal of Financial Economics*, 56(1): 125–149.
- Harris, Lawrence.** 1987. "Transaction Data Tests of the Mixture of Distributions Hypothesis." *Journal of Financial and Quantitative Analysis*, 22(2): 127–141.
- Harris, Lawrence.** 1994. "Minimum Price Variations, Discrete Bid-Ask Spreads, and Quotation Sizes." *Review of Financial Studies*, 7(1): 149–178.
- Hasbrouck, Joel.** 1999. "Trading Fast and Slow: Security Market Events in Real Time." Working Paper, New York University.
- Hasbrouck, Joel, George Sofianos, and Deborah Sosebee.** 1993. "New York Stock Exchange Systems and Trading Procedures." NYSE Working Paper 93-01.
- Hendershott, Terrence, Charles Jones, and Avanidhar Subrahmanyam.** 2011. "Does Algorithmic Trading Improve Liquidity?" *Journal of Finance*, 66(1): 1–33.
- Jones, Charles M., Gautam Kaul, and Marc L. Lipson.** 1994. "Information, Trading, and Volatility." *Journal of Financial Economics*, 36(1): 127–154.
- Keim, Donald B., and Ananth Madhavan.** 1995. "Anatomy of the Trading Process: Empirical Evidence on the Behavior of Institutional Trades." *Journal of Financial Economics*, 37(3): 371–398.
- Kirilenko, Andrei, Albert S. Kyle, Mehrdad Samadi, and Tugkan Tuzun.** 2015. "The Flash Crash: The Impact of High Frequency Trading on an Electronic Market." Working Paper, University of Maryland.
- Kyle, Albert S., and Anna A. Obizhaeva.** 2016. "Market Microstructure Invariance: Empirical Hypotheses." *Econometrica*, accepted for publication, available at <http://dx.doi.org/10.2139/ssrn.2722524>.
- Moulton, Pamela.** 2005. "You Can't Always Get What You Want: Trade-Size Clustering and Quantity Choice in Liquidity." *Journal of Financial Economics*, 78(5): 89–119.
- Obizhaeva, Anna A., and Jiang Wang.** 2013. "Optimal Trading Strategy and Supply/Demand Dynamics." *Journal of Financial Markets*, 16(1): 1–32.
- O'Hara, Maureen, Chen Yao, and Mao Ye.** 2012. "What Is Not There: The Odd-Lot Bias in TAQ Data." Working Paper.
- Schultz, Paul.** 2000. "Stock Splits, Tick Size, and Sponsorship." *Journal of Finance*, 55(1): 429–450.
- Smith, Jeffrey W., James P. Selway, and Timothy D. McCormick.** 1998. "The Nasdaq Stock Market: Historical Background and Current Operation." NASD Working Paper 98-01.

Tauchen, George E., and Mark Pitts. 1983. "The Price Variability-Volume Relationship on Speculative Markets." *Journal of Business and Economic Statistics*, 51: 485–505.

TABLE 1—DESCRIPTIVE STATISTICS

Volume Groups:	All	1	2	3	4	5	6	7	8	9	10
<i>Panel A: 1993–2000</i>											
Avg. Print Size (\$)	23,629	11,441	27,378	36,432	43,558	49,274	53,421	60,364	67,904	78,139	89,338
Med. (VW) Print Size (\$)	111,195	48,999	158,303	190,228	218,529	218,038	250,510	255,284	291,867	317,847	373,248
Med. (TW) Print Size (\$)	9,363	5,682	10,842	13,407	15,402	16,910	18,024	20,049	21,909	24,599	28,567
Avg # of Prints, γ	142	17	73	126	185	257	333	409	549	852	2,830
Avg. Daily Volume (\$ 1,000)	6,186	151	1,197	2,808	5,051	7,999	11,263	15,847	23,901	42,408	176,985
Avg. Daily Volatility	0.040	0.045	0.031	0.032	0.030	0.030	0.030	0.029	0.029	0.029	0.029
Avg. Price (\$)	17.58	10.25	20.38	24.50	27.84	31.57	34.73	38.34	43.03	49.98	64.43
100-Shares: % Prints/ %Vol	16/2	15/2	17/2	18/2	19/2	20/2	21/2	21/2	21/2	22/2	25/3
1,000-Shares: % Prints/ %Vol	18/14	18/15	18/13	17/12	16/12	15/11	15/11	14/10	14/10	13/10	13/10
Even Lots: %Prints/ % Vol	80/61	80/63	80/60	80/59	80/58	80/57	80/57	79/56	79/55	79/56	81/58
# Obs	634,322	391,611	93,732	37,272	33,145	14,896	14,155	13,039	12,276	11,831	12,365
<i>Panel B: 2001–14</i>											
Avg. Print Size (\$)	6,424	3,645	6,715	8,776	10,379	11,913	13,308	14,863	17,003	19,852	26,335
Med. (VW) Print Size (\$)	29,900	27,916	23,464	26,460	29,278	27,822	41,432	36,663	42,728	53,538	83,227
Med. (TW) Print Size (\$)	2,895	1,642	3,244	4,163	4,793	5,419	6,005	6,647	7,469	8,378	10,533
Avg # of Prints, γ	3,013	399	2,178	3,584	5,069	6,733	8,323	10,062	12,846	17,651	37,495
Avg. Daily Volume (\$ 1,000)	24,408	937	8,295	16,961	27,429	40,200	54,090	72,872	102,484	162,619	503,229
Avg. Daily Volatility	0.031	0.035	0.026	0.025	0.024	0.024	0.023	0.022	0.022	0.022	0.022
Avg. Price (\$)	21.16	12.56	26.77	32.00	34.92	37.81	40.92	44.46	48.24	51.77	63.33
100-Shares: % Prints/ %Vol	57/25	56/23	64/33	61/31	60/29	58/28	57/28	56/28	55/27	53/25	49/20
1,000-Shares: % Prints/ %Vol	3/6	4/6	2/4	2/4	2/4	2/4	2/4	2/4	2/4	3/4	3/5
Even Lots: %Prints/ % Vol	86/63	85/60	90/68	89/68	89/67	88/66	88/65	87/65	87/64	86/62	84/59
# Obs	749,535	477,127	97,347	39,485	36,846	17,455	16,780	16,222	15,858	15,815	16,600

This table reports descriptive statistics for securities and prints. Each observation represents averages for one security over one month. Panel A reports statistics for data from February 1993 to December 2000. Panel B reports statistics for data from January 2001 to December 2014. Both panels show the average print size, the trade-weighted median print size, the volume-weighted median print size (in dollars), the average number of prints per day, the daily dollar volume (in thousands of dollars), the average volatility of daily returns, the average price, and the percentage of trades and the percentage of volume in the 100-share lot, in the 1,000-share lot, and in the even lots for all samples as well as for 10 volume groups. Volume groups are based on average dollar trading volume with thresholds corresponding to the 30th, 50th, 60th, 70th, 75th, 80th, 85th, 90th, and 95th percentiles of the dollar volume for NYSE-listed common stocks. Volume group 1 has stocks with the lowest volume, and volume group 10 has stocks with the highest volume.

TABLE 2—OLS ESTIMATES OF NUMBER OF TAQ PRINTS.

	All Stocks		NYSE/AMEX		Nasdaq	
	1993–2000	2001–14	1993–2000	2001–14	1993–2000	2001–14
$\mu_{n,0}$	6.147	8.513	6.109	8.396	6.143	8.646
	0.017	0.041	0.012	0.040	0.020	0.048
$\mu_{n,t}$	0.093	0.148	0.042	0.178	0.141	0.115
	0.007	0.013	0.005	0.012	0.008	0.015
$a_{n,0}$	0.666	0.790	0.626	0.760	0.679	0.816
	0.002	0.005	0.001	0.005	0.002	0.006
$a_{n,t}$	-0.001	0.007	0.002	0.006	0.003	0.003
	0.001	0.001	0.000	0.001	0.001	0.002
Adj- R^2	0.87	0.92	0.91	0.93	0.86	0.91
# of Obs	6,621	4,452	2,189	1,759	4,432	2,694

This table presents Fama-MacBeth estimates μ_n and a_n from monthly regressions

$$\ln [\bar{N}_i] = \mu_n + a_n \cdot \ln \left[\frac{W_i}{W_*} \right] + \tilde{\epsilon}_i.$$

For each month, there is one observation for each stock i . The value of \bar{N}_i is the average number of prints per day. Trading activity W_i is the product of average daily dollar volume $V_i \cdot P_i$ and the percentage standard deviation σ_i of daily returns in a given month. The scaling constant $W_* = (40)(10^6)(0.02)$ corresponds to the measure of trading activity for a benchmark stock with a price of \$40 per share, trading volume of 1 million shares per day, and daily volatility of 0.02. Newey-West standard errors are calculated with three lags relative to a linear time trend estimated by OLS regressions from the estimated coefficients $\hat{\mu}_{n,T}$ and $\hat{a}_{n,T}$ for each month: $\hat{\mu}_{n,T} = \mu_{n,0} + \mu_{n,t} \cdot (T - \bar{T})/12 + \tilde{\epsilon}_{\bar{T}}$ and $\hat{a}_{n,T} = a_{n,0} + a_{n,t} \cdot (T - \bar{T})/12 + \tilde{\epsilon}_{\bar{T}}$, where T is the number of months from the beginning of the sample and \bar{T} is the mean month. “Adj- R^2 ” denotes the adjusted R^2 averaged over monthly regressions, and “# of Obs” denotes the number of stocks averaged over monthly regressions. The estimates are reported for 1993–2000 and 2001–14 subperiods.

TABLE 3—REGRESSION ESTIMATES OF TAQ PRINT SIZES, FEBRUARY 1993 TO DECEMBER 2000

	Trade-Weighted Distribution				Volume-Weighted Distribution			
	Mean	20th	50th	80th	Mean	20th	50th	80th
$\mu_{x,0}$	-7.238	-8.495	-7.289	-6.260	-4.684	-6.364	-4.887	-3.326
	0.020	0.027	0.031	0.014	0.020	0.016	0.019	0.026
$\mu_{x,t}$	-0.047	-0.039	-0.046	-0.064	-0.137	-0.101	-0.157	-0.153
	0.008	0.011	0.012	0.006	0.008	0.006	0.008	0.011
$a_{x,0}$	-0.741	-0.781	-0.750	-0.725	-0.560	-0.661	-0.579	-0.481
	0.002	0.003	0.002	0.003	0.003	0.002	0.003	0.003
$a_{x,t}$	0.008	0.012	0.007	0.005	-0.007	-0.001	-0.009	-0.009
	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001
Adj- R^2	0.90	0.87	0.88	0.88	0.69	0.83	0.68	0.52
# of Obs	6,621	6,621	6,621	6,621	6,621	6,621	6,621	6,621

This table presents Fama-MacBeth estimates μ_x and a_x from the monthly regressions of the mean and percentiles of print size on trading activity W for the sample from February 1993 to December 2000. The coefficients μ_Q and a_Q are based on monthly regressions

$$\ln \left[\frac{|X_i|}{V_i} \right] = \mu_x + a_x \cdot \ln \left[\frac{W_i}{W_*} \right] + \tilde{\epsilon}_i,$$

where the left-hand side is either the mean or the p th (20th, 50th and 80th) percentile of the distribution of logarithms of (unsigned) print sizes $|X_i|$, expressed as a fraction of daily volume V_i in a given month. The means and percentiles are calculated based on both trade-weighted and volume-weighted distributions. For each month, there is one observation for each stock i , with trading activity W_i defined as the product of the average daily dollar volume $V_i \cdot P_i$ and the percentage standard deviation σ_i of daily returns. The scaling constant $W_* = (40)(10^6)(0.02)$ corresponds to the trading activity of the benchmark stock with a price of \$40 per share, trading volume of 1 million shares per day, and volatility of 2 percent per day. Newey-West standard errors are calculated with three lags relative to a linear time trend estimated by OLS regressions from the estimated coefficients $\hat{\mu}_{x,T}$ and $\hat{a}_{x,T}$ for each month: $\hat{\mu}_{x,T} = \mu_{x,0} + \mu_{x,t} \cdot (T - \bar{T})/12 + \tilde{\epsilon}_T$ and $\hat{a}_{x,T} = a_{x,0} + a_{x,t} \cdot (T - \bar{T})/12 + \tilde{\epsilon}_T$, where T is the number of months from the beginning of the sample and \bar{T} is the mean month. “Adj- R^2 ” denotes the adjusted R^2 averaged over monthly regressions, and “# of Obs” denotes the number of stocks averaged over monthly regressions.

TABLE 4—REGRESSION ESTIMATES OF TAQ PRINT SIZES, JANUARY 2001 TO DECEMBER 2014

	Trade-Weighted Distribution				Volume-Weighted Distribution			
	Mean	20th	50th	80th	Mean	20th	50th	80th
$\mu_{x,0}$	-9.029	-9.519	-9.268	-8.633	-7.420	-9.004	-8.123	-6.304
	0.030	0.024	0.038	0.039	0.058	0.047	0.059	0.059
$\mu_{x,t}$	-0.097	-0.013	-0.084	-0.173	-0.192	-0.091	-0.251	-0.251
	0.009	0.007	0.012	0.012	0.018	0.015	0.018	0.018
$a_{x,0}$	-0.779	-0.757	-0.769	-0.811	-0.776	-0.796	-0.843	-0.779
	0.003	0.003	0.003	0.004	0.007	0.006	0.009	0.010
$a_{x,t}$	0.000	0.005	0.003	-0.006	-0.018	-0.005	-0.021	-0.026
	0.001	0.001	0.001	0.001	0.002	0.002	0.002	0.003
Adj- R^2	0.907	0.866	0.884	0.913	0.892	0.892	0.862	0.760
# of Obs	4,452	4,452	4,452	4,452	4,452	4,452	4,452	4,452

This table presents Fama-MacBeth estimates μ_x and a_x from the monthly regressions of the mean and percentiles of print size on trading activity W for the sample from January 2001 to December 2014. The coefficients μ_x and a_x are based on monthly regressions

$$\ln \left[\frac{|X_i|}{V_i} \right] = \mu_x + a_x \cdot \ln \left[\frac{W_i}{W_*} \right] + \tilde{\epsilon}_i,$$

where the left-hand side is either the mean or the p th (20th, 50th and 80th) percentile of the distribution of logarithms of (unsigned) print sizes $|X_i|$, expressed as a fraction of daily volume V_i in a given month. The means and percentiles are calculated based on both trade- and volume-weighted distributions. For each month, there is one observation for each stock i , with trading activity W_i defined as the product of the average daily dollar volume $V_i \cdot P_i$ and the percentage standard deviation σ_i of daily returns. The scaling constant $W_* = (40)(10^6)(0.02)$ corresponds to the trading activity of the benchmark stock with a price of \$40 per share, trading volume of 1 million shares per day, and volatility of 2 percent per day. Newey-West standard errors are calculated with three lags relative to a linear time trend estimated by OLS regressions from the estimated coefficients $\hat{\mu}_{x,T}$ and $\hat{a}_{x,T}$ for each month: $\hat{\mu}_{x,T} = \mu_{x,0} + \mu_{x,t} \cdot (T - \bar{T})/12 + \tilde{\epsilon}_T$ and $\hat{a}_{x,T} = a_{x,0} + a_{x,t} \cdot (T - \bar{T})/12 + \tilde{\epsilon}_T$, where T is the number of months from the beginning of the sample and \bar{T} is the mean month. “Adj- R^2 ” denotes the adjusted R^2 averaged over monthly regressions, and “# of Obs” denotes the number of stocks averaged over monthly regressions.

TABLE 5—OLS ESTIMATES OF NUMBER OF TAQ PRINTS WITH EFFECTIVE VOLATILITY.

	All Stocks		NYSE/AMEX		Nasdaq	
	1993–2000	2001–14	1993–2000	2001–14	1993–2000	2001–14
$\mu_{n,0}$	6.270	7.952	6.316	7.998	6.247	7.942
	0.016	0.023	0.009	0.025	0.020	0.027
$\mu_{n,t}$	0.087	0.116	0.024	0.151	0.125	0.093
	0.007	0.007	0.004	0.007	0.009	0.008
$a_{\sigma,0}$	-0.471	-0.608	-0.338	-0.475	-0.497	-0.676
	0.003	0.009	0.005	0.008	0.003	0.010
$a_{\sigma,t}$	-0.007	-0.020	-0.007	-0.027	-0.023	-0.013
	0.001	0.003	0.002	0.002	0.001	0.003
Adj- R^2	0.918	0.955	0.924	0.955	0.917	0.959
# of Obs	6,621	4,452	2,189	1,759	4,432	2,694
Regression with Coefficient on Effective Price Volatility $a_{\sigma} = 0$.						
Adj- R^2	0.873	0.899	0.908	0.917	0.857	0.881
Regression with Separate Coefficients for Price, Volume, and Volatility.						
Adj- R^2	0.928	0.970	0.940	0.974	0.923	0.975

This table presents Fama-MacBeth estimates μ_n and a_{σ} from monthly regressions

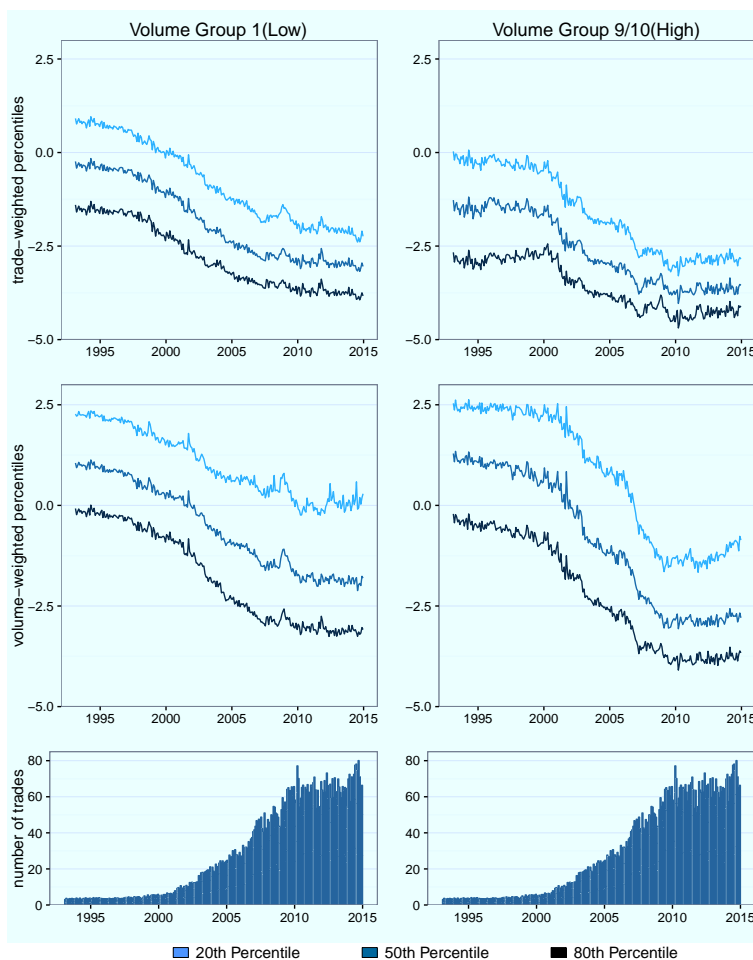
$$\ln [\bar{N}_i] = \mu_n + \frac{2}{3} \cdot \ln \left[\frac{W_i}{W_*} \right] + a_{\sigma} \cdot \ln \left[\frac{P_i \cdot \sigma_i}{P_* \cdot \sigma_*} \cdot \left(\frac{W_i}{W_*} \right)^{-1/3} \right] + \tilde{\epsilon}_i.$$

For each month, there is one observation for each stock i , with trading activity W_i defined as the product of the average daily dollar volume $V_i \cdot P_i$ and the percentage standard deviation σ_i of daily returns. Effective price volatility is defined as $P_i \cdot \sigma_i \cdot \left(\frac{W_i}{W_*} \right)^{-1/3}$, with the effective price volatility of the benchmark stocks $P_* \cdot \sigma_*$ equal to $40 \cdot 0.02$. The value of \bar{N}_i is the average number of prints per day. The scaling constant $W_* = (40)(10^6)(0.02)$ corresponds to the measure of trading activity for the benchmark stock with a price of \$40 per share, trading volume of 1 million shares per day, and daily volatility of 2 percent. Newey-West standard errors are calculated with three lags relative to a linear time trend estimated by OLS regressions from the estimated coefficients $\hat{\mu}_{n,T}$ and $\hat{a}_{\sigma,T}$ for each month: $\hat{\mu}_{n,T} = \mu_{n,0} + \mu_{n,t} \cdot (T - \bar{T})/12 + \tilde{\epsilon}_{\hat{\mu}}$ and $\hat{a}_{\sigma,T} = a_{\sigma,0} + a_{\sigma,t} \cdot (T - \bar{T})/12 + \tilde{\epsilon}_{\hat{a}}$, where T is the number of months from the beginning of the sample and \bar{T} is the mean month. “Adj- R^2 ” denotes the adjusted R^2 averaged over monthly regressions. The table also reports the average R^2 from the restricted regressions with $a_{\sigma} = 0$ as well as the average R^2 from unconstrained regressions

$$(19) \quad \ln [\bar{N}] = \mu_n + \frac{2}{3} \ln \left[\frac{W_i}{W_*} \right] + b_1 \cdot \ln \left[\frac{V_i}{(10^6)} \right] + b_2 \cdot \ln \left[\frac{P_i}{(40)} \right] + b_3 \cdot \ln \left[\frac{\sigma_i}{(0.02)} \right] + \tilde{\epsilon}.$$

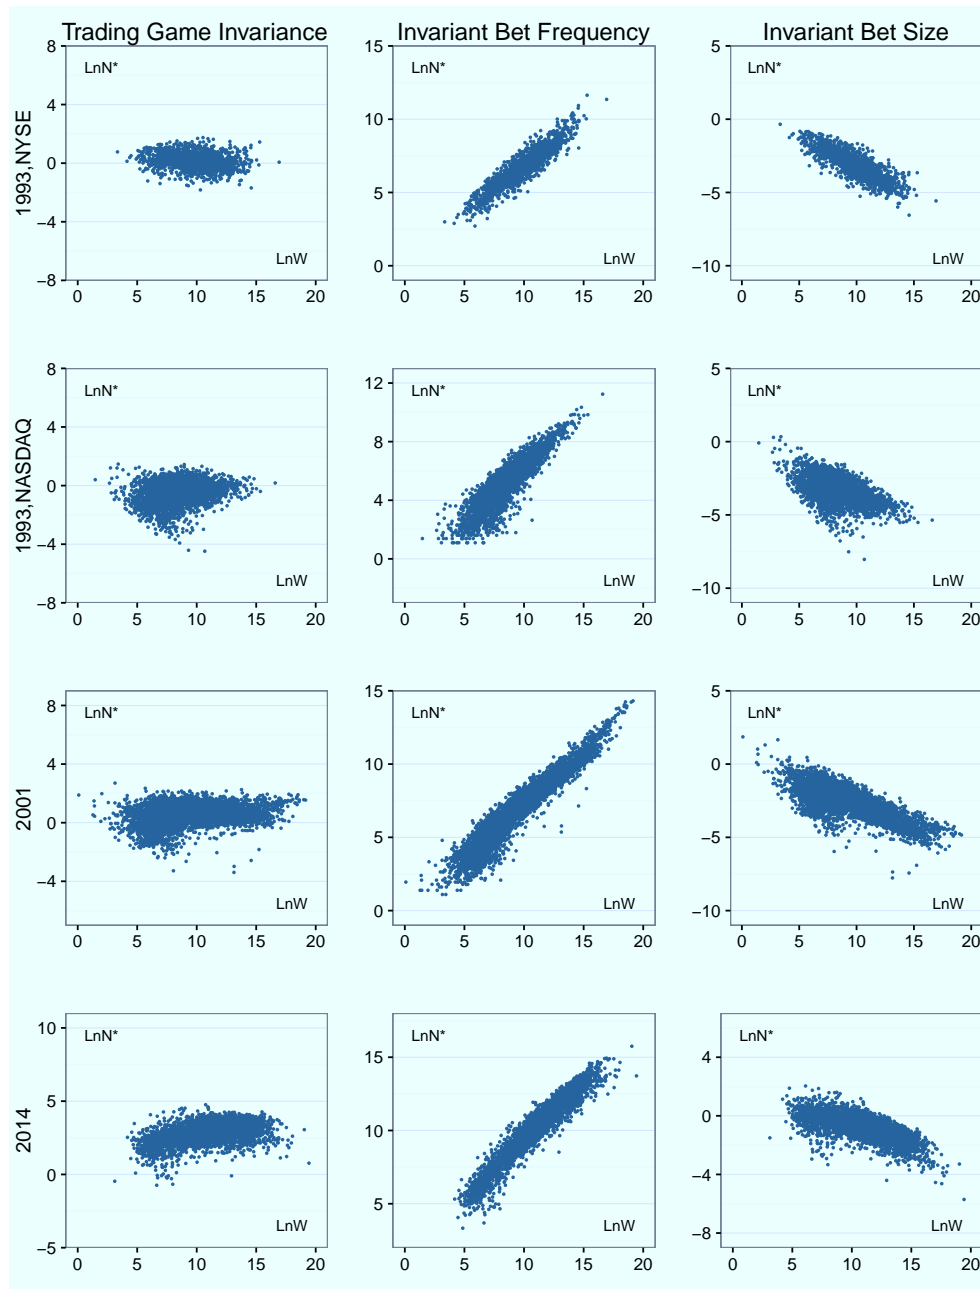
“# of Obs” is the number of stocks averaged over monthly regressions. The estimates are reported for 1993–2000 and 2001–14 subperiods.

FIGURE 1. TIME SERIES OF PERCENTILES OF SCALED TAQ PRINT SIZE AND MEAN NUMBER OF PRINTS, 1993–2014.



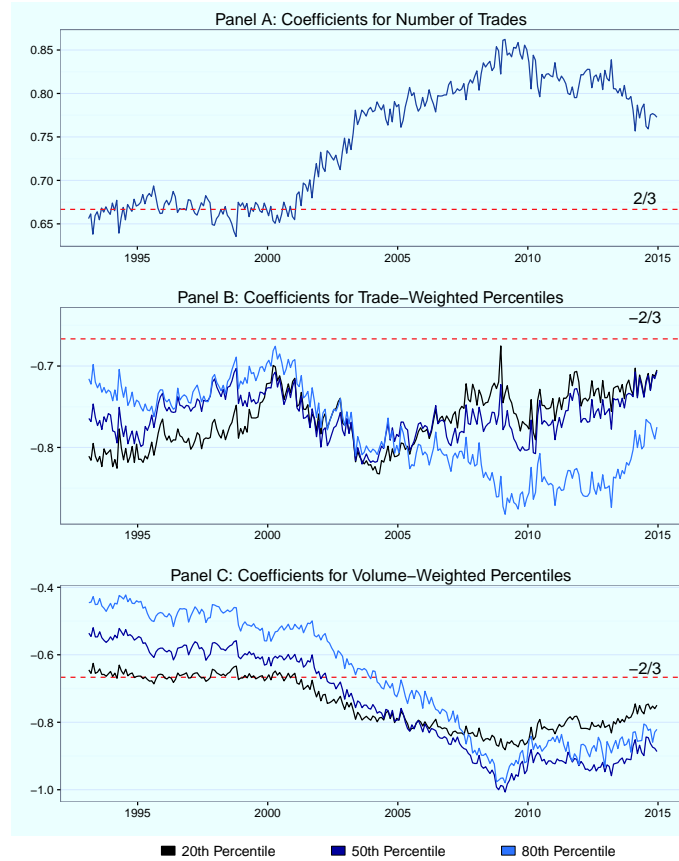
The figure shows the dynamics of the 20th, 50th, and 80th percentiles for logarithms of the pooled scaled print sizes as well as the means of the scaled number of prints per month from 1993 to 2014. Volume groups are based on average dollar trading volume with thresholds corresponding to the 30th, 50th, 60th, 70th, 75th, 80th, 85th, 90th, and 95th percentiles of the dollar volume for common NYSE-listed stocks. Trade-weighted percentiles and volume-weighted percentiles are shown for stocks in volume group 1 (low volume) and volume groups 9 and 10 (high volume). For each print, the logarithm of scaled print size is calculated based on the midpoint of the print size bin, scaled according to the model of trading game invariance—that is, $\ln(W_i^{2/3} \cdot |X_i|/V_i)$, where $|X_i|$ is a midpoint of a print size bin in shares, V_i is the average daily volume in shares, and W_i is the measure of trading activity equal to the product of dollar volume and returns standard deviation. The scaled number of prints per month is calculated as $\bar{N}_{m,i} \cdot W_i^{-2/3}$, where $\bar{N}_{m,i}$ is the number of trades per month. The stock-level distributions of scaled print sizes are averaged across stocks for volume groups 1 and 9–10 in a given month. The trade- and volume-weighted percentiles are plotted on this figure.

FIGURE 2. THE SCALED NUMBER OF TAQ PRINTS RELATIVE TO TRADING ACTIVITY FOR THREE MODELS.



The figure shows the logarithm of the scaled number of prints across different levels of the logarithm of trading activity W_i . The scaled number of prints is defined by \bar{N}_i/W_i^α , with $\alpha = 2/3$ for the model of trading game invariance, $\alpha = 0$ for the model of invariant bet frequency, and $\alpha = 1$ for the model of invariant bet size. Four subsamples are considered: NYSE-listed stocks in April 1993, Nasdaq-listed stocks in April 1993, both NYSE and Nasdaq stocks in April 2001 and both Nasdaq and NYSE stocks in April 2014. Trading activity W_i is calculated as the product of average daily dollar $P_i \cdot V_i$ volume and the percentage standard deviation of daily returns σ_i for a given month.

FIGURE 3. TIME SERIES OF MONTHLY OLS COEFFICIENT ESTIMATES FOR NUMBER OF TRADES, TRADE-WEIGHTED PERCENTILES, AND VOLUME-WEIGHTED PERCENTILES, 1993–2014.



The figure shows the dynamics of coefficients from regressions of number of prints and various percentiles on the measure of trading activity W_i from 1993 to 2014. Panel A shows the coefficient a_n from monthly regressions


$$\ln [\bar{N}_i] = \mu_n + a_n \cdot \ln \left[\frac{W_i}{W_*} \right] + \tilde{\epsilon}_i,$$

where \bar{N}_i is the average number of prints per day in a given month. The model of trading game invariance predicts $a_n = 2/3$, and alternative models predict that $a_n = 0$ or $a_n = 1$. Panel B shows the coefficient a_x from monthly regressions

$$\ln \left[\frac{\tilde{X}_i}{V_i} \right] = \mu_x + a_x \cdot \ln \left[\frac{W_i}{W_*} \right] + \tilde{\epsilon}_i,$$

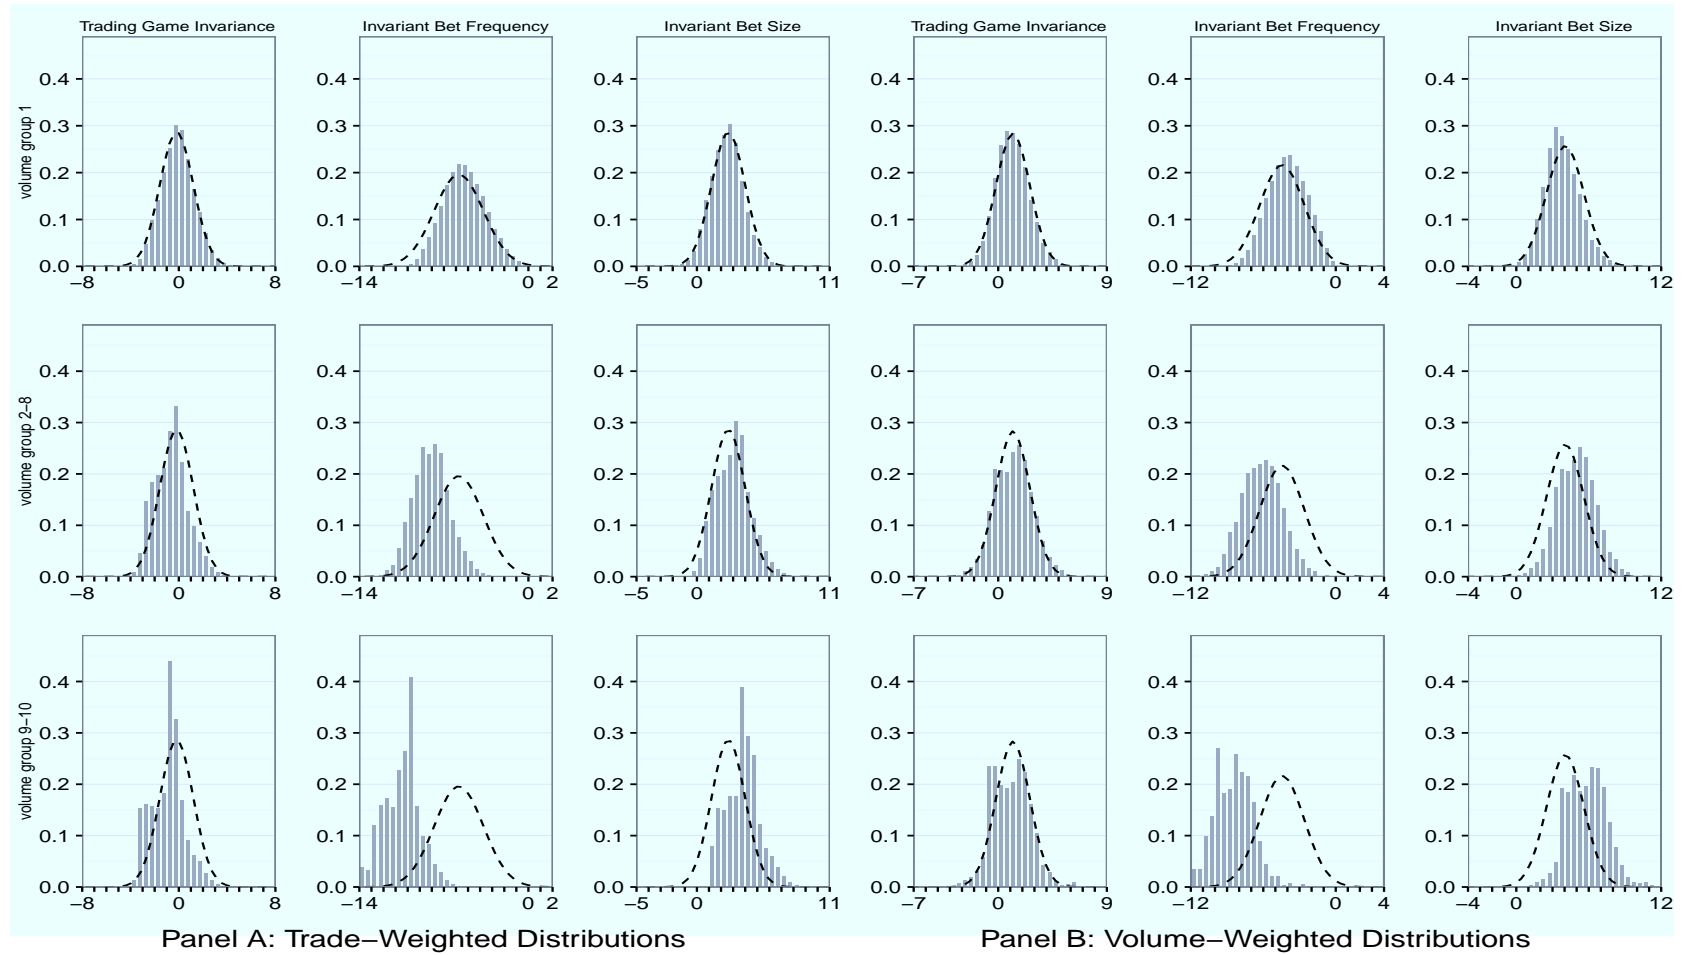
where the left-hand side is the p th (20th, 50th and 80th) percentiles of the distribution of logarithms of print sizes \tilde{X}_i . The model of trading game invariance predicts $a_x = -2/3$, and alternative models predict that $a_x = 0$ or $a_x = -1$. Panel C shows the coefficient a_x from similar monthly regressions, but these regressions are based on percentiles Q_i^p , where percentiles are calculated based on the contribution to total trading volume. The model of trading game invariance predicts $a_x = -2/3$, and alternative models predict that $a_x = 0$ or $a_x = -1$. Trading activity W_i is defined as the product of dollar volume and daily percentage standard deviation of returns, and W_* measures trading activity of the benchmark stock.

FIGURE 4. TRADE-WEIGHTED AND VOLUME-WEIGHTED DISTRIBUTIONS OF SCALED TAQ PRINT SIZE FOR THREE MODELS, NYSE-LISTED STOCKS, APRIL 1993



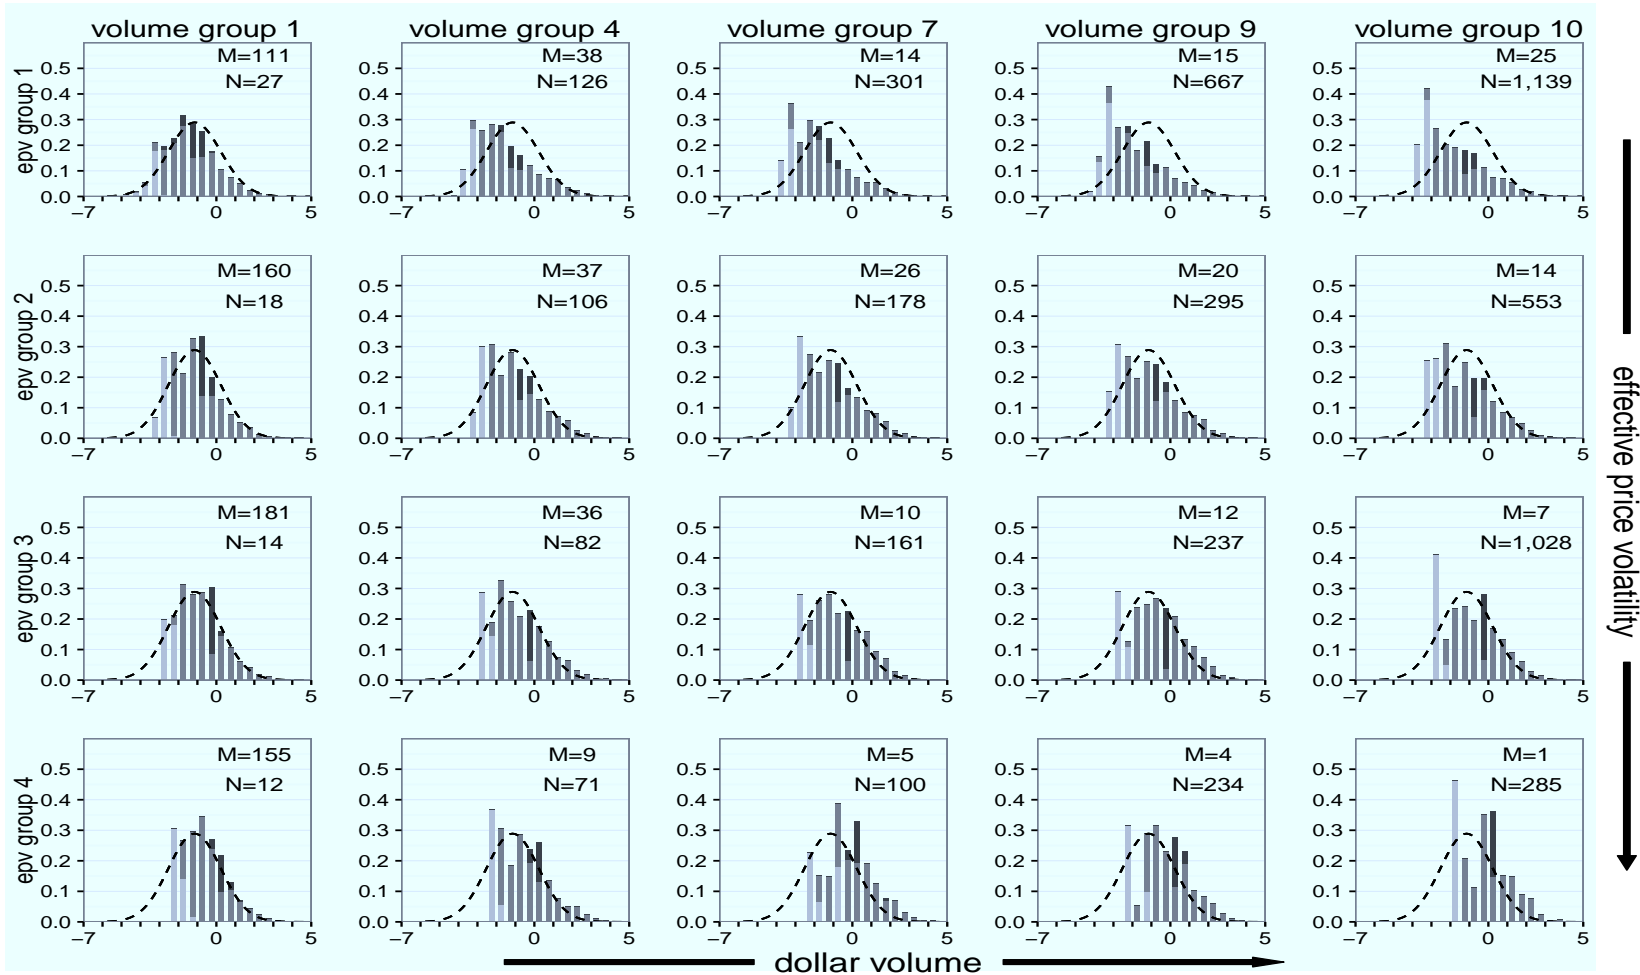
This figure shows the distribution of the logarithm of scaled print sizes for three different models for NYSE-listed stocks traded in April 1993. The print sizes are scaled as $W_i^\alpha \cdot |X_i|/V_i$, with $\alpha = 2/3$ for the model of invariant bet frequency, $\alpha = 0$ for the model of invariant bet frequency, and $\alpha = 1$ for the model of invariant bet size. Trading activity W_i is calculated as the product of dollar volume $P_i \cdot V_i$ and the daily percentage standard deviation of returns σ_i . Panel A shows trade-weighted distributions, and panel B shows volume-weighted distributions. The subplots show stock-level distributions averaged across stocks in volume group 1 (low volume), volume groups 2–8, and volume groups 9–10 (high volume). Volume groups are based on average dollar trading volume with thresholds corresponding to the 30th, 50th, 60th, 70th, 75th, 80th, 85th, 90th, and 95th percentiles of the dollar volume for NYSE-listed common stocks.

FIGURE 5. TRADE-WEIGHTED AND VOLUME-WEIGHTED DISTRIBUTIONS OF SCALED TAQ PRINT SIZE FOR THREE MODELS, NASDAQ-LISTED STOCKS, APRIL 1993



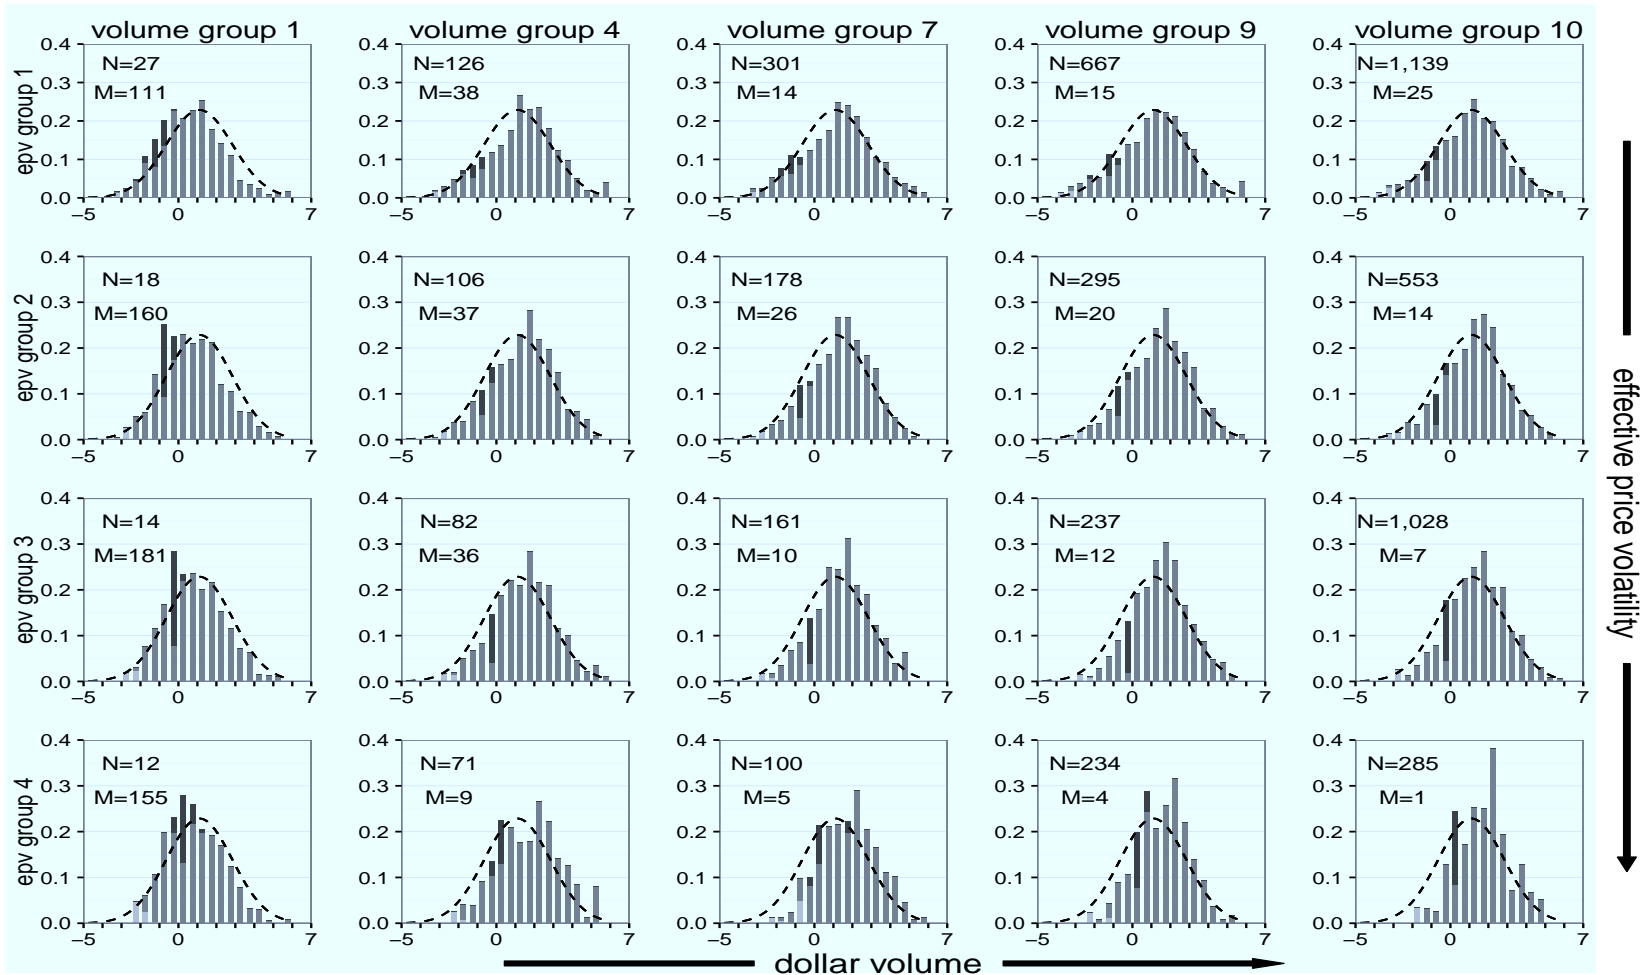
This figure shows the distribution of the logarithm of scaled print sizes for three different models for Nasdaq-listed stocks traded in April 1993. The print sizes are scaled as $W_i^\alpha \cdot |X_i|/V_i$, with $\alpha = 2/3$ for the model of invariant bet frequency, $\alpha = 0$ for the model of invariant bet frequency, and $\alpha = 1$ for the model of invariant bet size. Trading activity W_i is calculated as the product of dollar volume $P_i \cdot V_i$ and the daily percentage standard deviation of returns σ_i . Panel A shows trade-weighted distributions, and panel B shows volume-weighted distributions. The subplots show stock-level distributions averaged across stocks in volume group 1 (low volume), volume groups 2–9, and volume groups 9–10 (high volume). Volume groups are based on average dollar trading volume with thresholds corresponding to the 30th, 50th, 60th, 70th, 75th, 80th, 85th, 90th, and 95th percentiles of the dollar volume for Nasdaq-listed common stocks.

FIGURE 6. TRADE-WEIGHTED DISTRIBUTIONS OF SCALED TAQ PRINT SIZES, NYSE-LISTED STOCKS, APRIL 1993



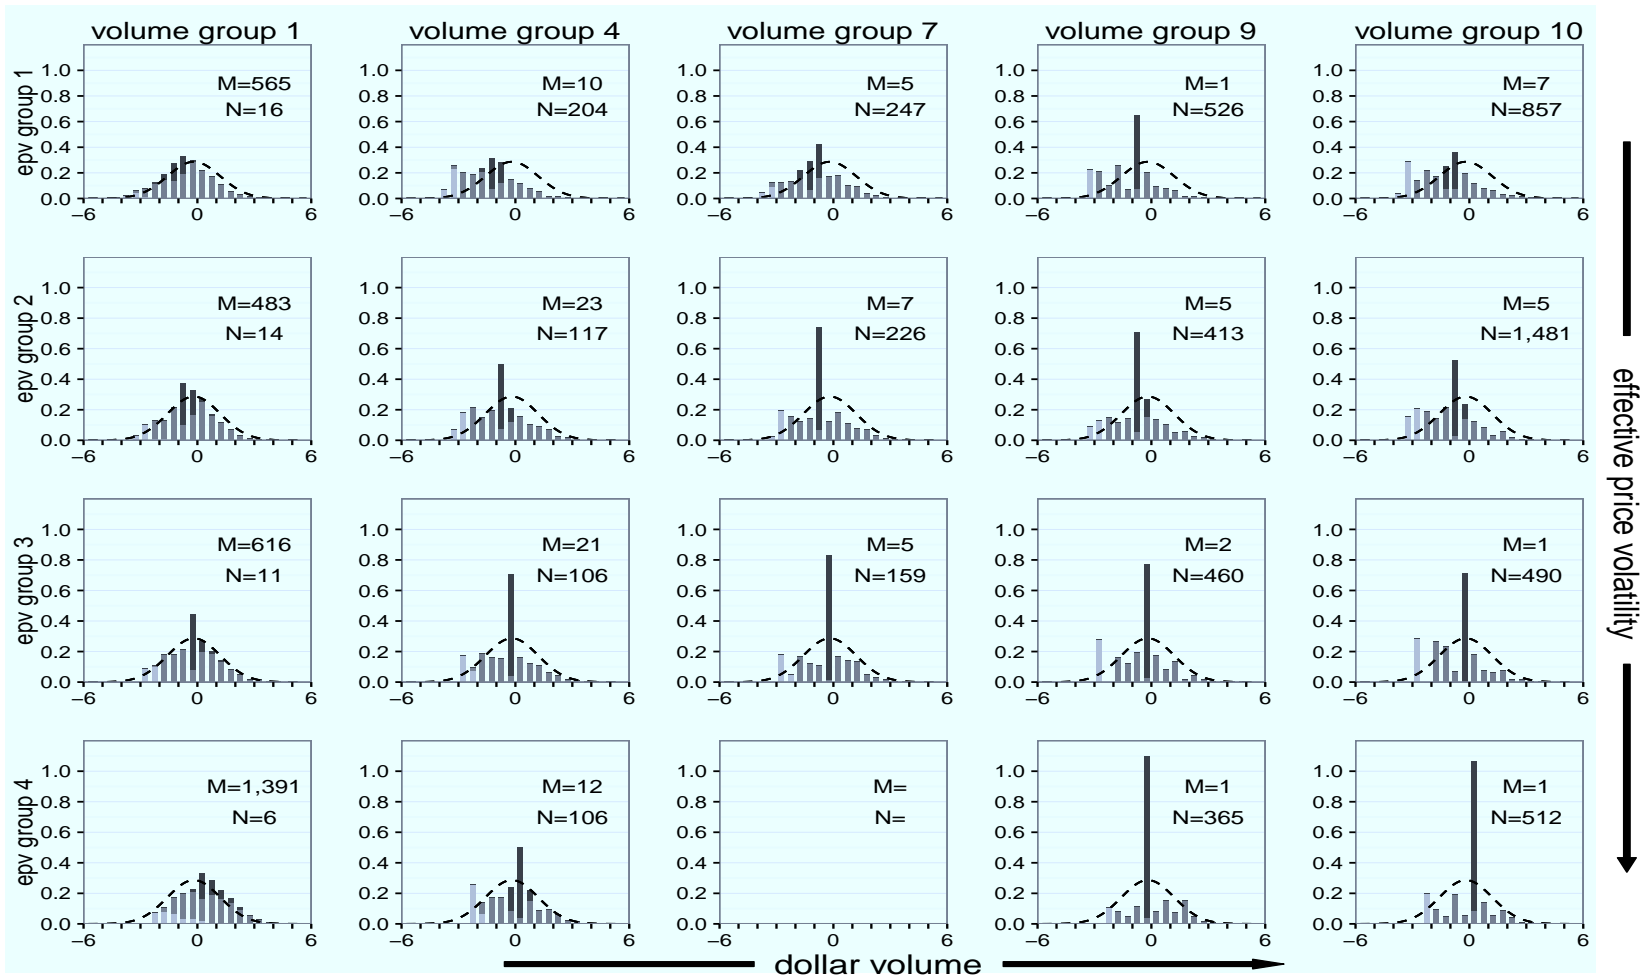
This figure shows distributions of the logarithms of scaled print sizes for NYSE stocks in April 1993. For each trade, the scaled print size is calculated as $\ln(W_i^{2/3} \cdot |X_i|/V_i)$ based on the invariance hypothesis, where $|X_i|$ is the midpoint of the print size bin in shares, V_i is the average daily volume in shares, and W_i measures trading activity as the product of dollar volume and the daily percentage standard deviation of returns. Ten volume groups are constructed based on average dollar trading volume with thresholds corresponding to the 30th, 50th, 60th, 70th, 75th, 80th, 85th, 90th, and 95th percentiles of the dollar volume for common NYSE-listed stocks. Four equally spaced volatility groups are constructed based on effective price volatility, defined as $P_i \cdot \sigma_i \cdot (W_i/W_*)^{-1/3}$. The subplots show stock-level distributions averaged across stocks for volume groups 1 (low volume), 4, 7, 9, and 10 (high volume) and for all four price volatility groups 1 (low price volatility), 2, 3, and 4 (high price volatility). The 100-share trades are highlighted in light gray; the 1,000-share trades are highlighted in dark gray. Each subplot also shows a normal distribution with the pooled average print size mean of -1.15 and standard deviation of 1.38. M is the number of stocks, and N is the average number of prints per day for the stocks in a given subgroup.

FIGURE 7. VOLUME-WEIGHTED DISTRIBUTIONS OF SCALED TAQ PRINT SIZES, NYSE-LISTED STOCKS, APRIL 1993



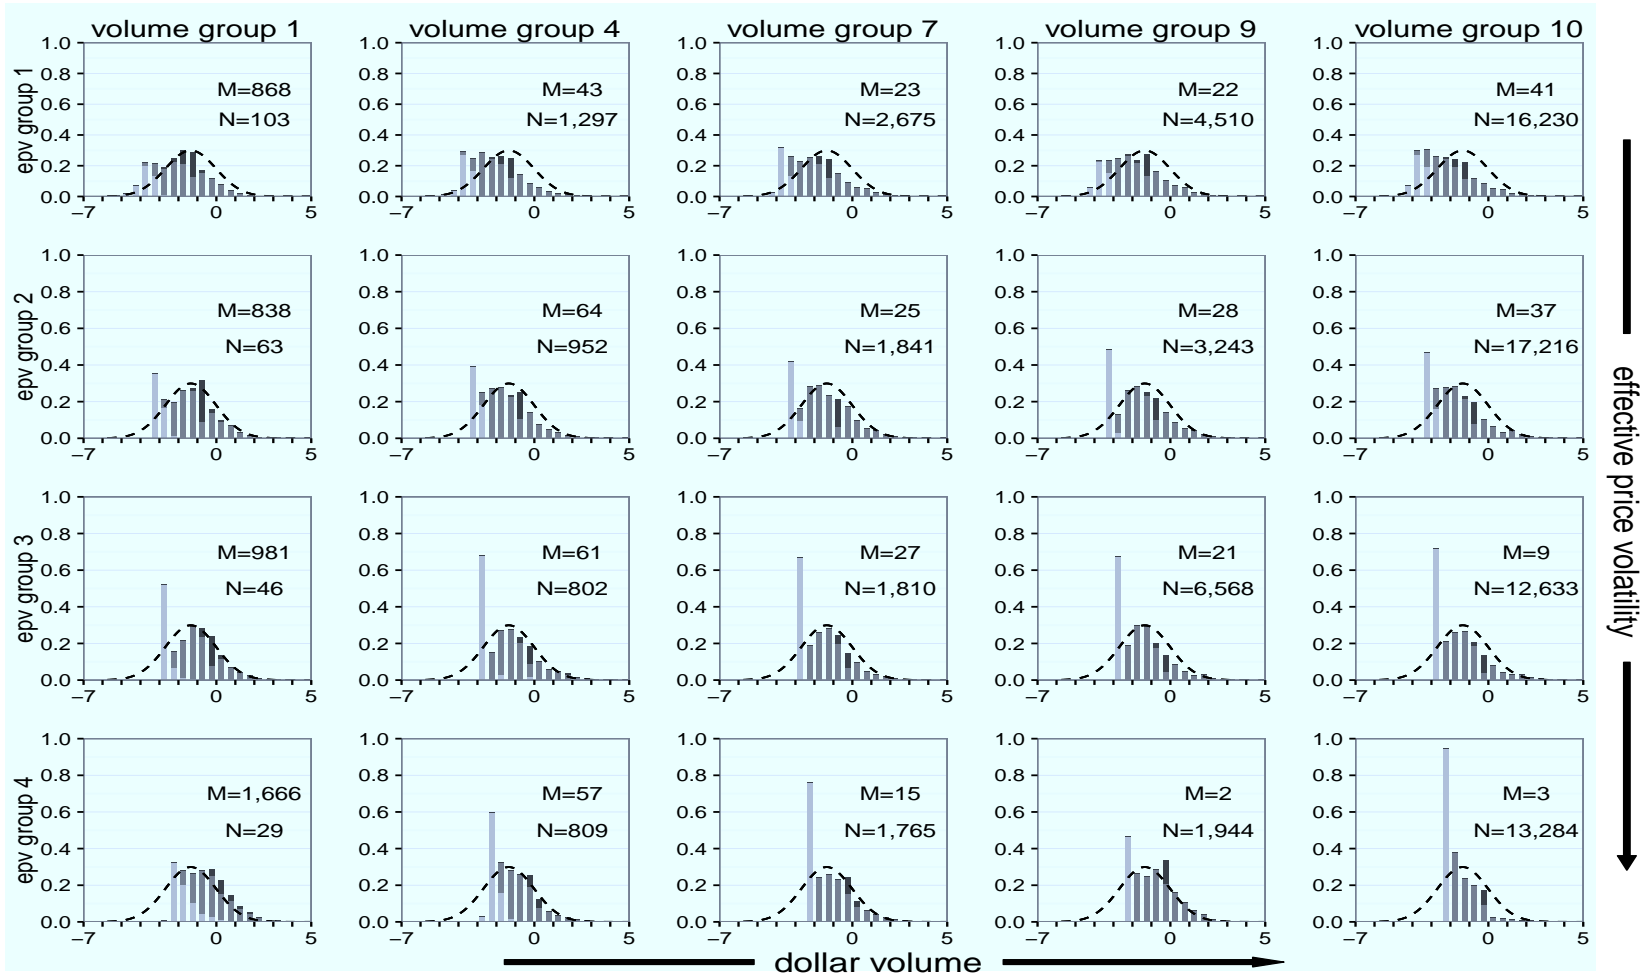
This figure shows distributions of total volume across different scaled print size bins for the NYSE stocks in April 1993. For each stock, the volume distribution is calculated as the contribution to the total volume by trades from a given trade size bin. The x-axis is the log of scaled print sizes, defined by $\ln(W_i^{2/3} \cdot \frac{|X_i|}{V_i})$ according to the invariance hypothesis, where $|X_i|$ is a print size in shares (midpoint of a bin), V_i is the average daily volume in shares, and W_i is the measure of trading activity equal to the product of dollar volume and returns standard deviation. Ten volume groups are constructed based on average dollar trading volume with thresholds corresponding to the 30th, 50th, 60th, 70th, 75th, 80th, 85th, 90th, and 95th percentiles of the dollar volume for NYSE-listed common stocks. Four equally spaced volatility groups are constructed based on effective price volatility, defined as $P_i \cdot \sigma_i \cdot (W_i/W_*)^{-1/3}$. The subplots show stock-level distributions averaged across stocks for volume groups 1 (low volume), 4, 7, 9, and 10 (high volume) and for all four price volatility groups 1 (low price volatility), 2, 3, and 4 (high price volatility). The 100-share trades are highlighted in light gray, and the 1,000-share trades are highlighted in dark gray. Each subplot also shows a normal distribution with the pooled average print size mean of 1.1 and standard deviation of 1.74. M is the number of stocks, and N is the average number of prints per day for the stocks in a given subgroup.

FIGURE 8. TRADE-WEIGHTED DISTRIBUTIONS OF SCALED TAQ PRINT SIZES, NASDAQ-LISTED STOCKS, APRIL 1993



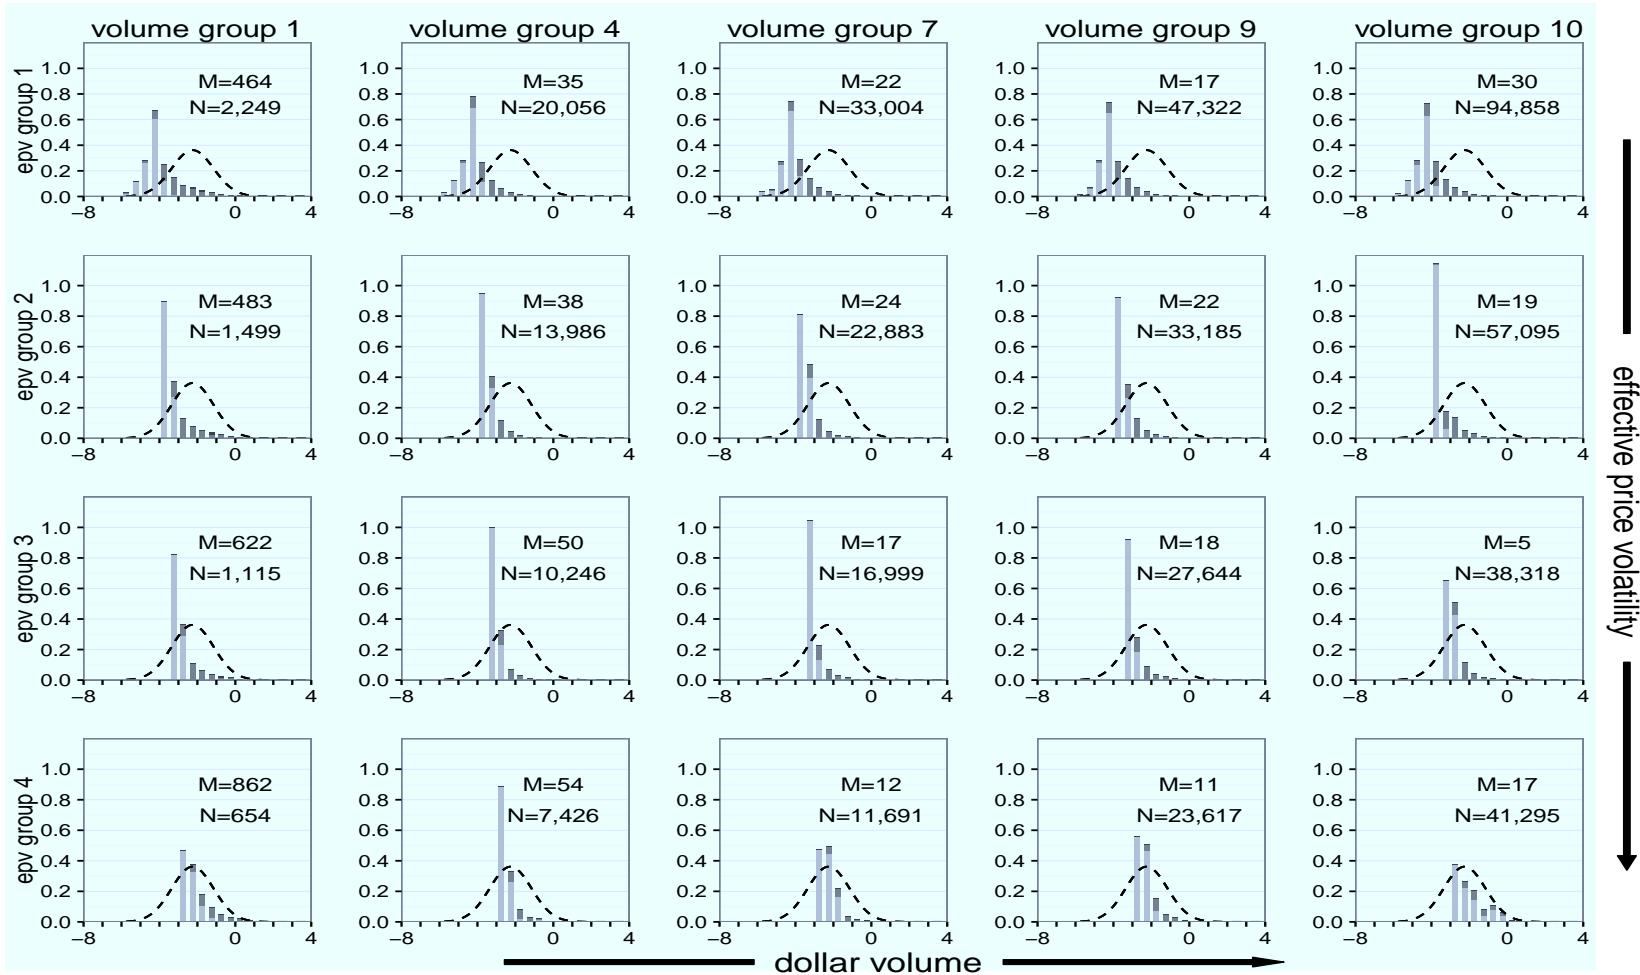
This figure shows distributions of the logarithms of scaled print sizes for Nasdaq stocks in April 1993. For each trade, the scaled print size is calculated as $\ln(W_i^{2/3} \cdot |X_i|/V_i)$ based on the invariance hypothesis, where $|X_i|$ is the midpoint of the print size bin in shares, V_i is the average daily volume in shares, and W_i measures trading activity as the product of dollar volume and the daily percentage standard deviation of returns. Ten volume groups are constructed based on average dollar trading volume with thresholds corresponding to the 30th, 50th, 60th, 70th, 75th, 80th, 85th, 90th, and 95th percentiles of the dollar volume for common NYSE-listed stocks. Four equally spaced volatility groups are constructed based on effective price volatility, defined as $P_i \cdot \sigma_i \cdot (W_i/W_*)^{-1/3}$. The subplots show stock-level distributions averaged across stocks for volume groups 1 (low volume), 4, 7, 9, and 10 (high volume) and for all four price volatility groups 1 (low price volatility), 2, 3, and 4 (high price volatility). The 100-share trades are highlighted in light gray; the 1,000-share trades are highlighted in dark gray. Each subplot also shows a normal distribution with the pooled average print size mean of -0.19 and standard deviation of 1.39. M is the number of stocks, and N is the average number of prints per day for the stocks in a given subgroup.

FIGURE 9. TRADE-WEIGHTED DISTRIBUTIONS OF SCALED TAQ PRINT SIZES, ALL STOCKS, APRIL 2001



This figure shows distributions of the logarithms of scaled print sizes for NYSE and Nasdaq stocks in April 2001. For each trade, the scaled print size is calculated as $\ln(W_i^{2/3} \cdot |X_i|/V_i)$ based on the invariance hypothesis, where $|X_i|$ is the midpoint of the print size bin in shares, V_i is the average daily volume in shares, and W_i measures trading activity as the product of dollar volume and the daily percentage standard deviation of returns. Ten volume groups are constructed based on average dollar trading volume with thresholds corresponding to the 30th, 50th, 60th, 70th, 75th, 80th, 85th, 90th, and 95th percentiles of the dollar volume for common NYSE-listed stocks. Four equally spaced volatility groups are constructed based on effective price volatility, defined as $P_i \cdot \sigma_i \cdot (W_i/W_*)^{-1/3}$. The subplots show stock-level distributions averaged across stocks for volume groups 1 (low volume), 4, 7, 9, and 10 (high volume) and for all four price volatility groups 1 (low price volatility), 2, 3, and 4 (high price volatility). The 100-share trades are highlighted in light gray; the 1,000-share trades are highlighted in dark gray. Each subplot also shows a normal distribution with the pooled average print size mean of -1.35 and standard deviation of 1.33 . M is the number of stocks, and N is the average number of prints per day for the stocks in a given subgroup.

FIGURE 10. TRADE-WEIGHTED DISTRIBUTIONS OF SCALED TAQ PRINT SIZES, ALL STOCKS, APRIL 2014



This figure shows distributions of the logarithms of scaled print sizes for NYSE and Nasdaq stocks in April 2014. For each trade, the scaled print size is calculated as $\ln(W_i^{2/3} \cdot |X_i|/V_i)$ based on the invariance hypothesis, where $|X_i|$ is the midpoint of the print size bin in shares, V_i is the average daily volume in shares, and W_i measures trading activity as the product of dollar volume and the daily percentage standard deviation of returns. Ten volume groups are constructed based on average dollar trading volume with thresholds corresponding to the 30th, 50th, 60th, 70th, 75th, 80th, 85th, 90th, and 95th percentiles of the dollar volume for common NYSE-listed stocks. Four equally spaced volatility groups are constructed based on effective price volatility, defined as $P_i \cdot \sigma_i \cdot (W_i/W_*)^{-1/3}$. The subplots show stock-level distributions averaged across stocks for volume groups 1 (low volume), 4, 7, 9, and 10 (high volume) and for all four price volatility groups 1 (low price volatility), 2, 3, and 4 (high price volatility). The 100-share trades are highlighted in light gray; the 1,000-share trades are highlighted in dark gray. Each subplot also shows a normal distribution with the pooled average print size mean of -2.25 and standard deviation of 1.10 . M is the number of stocks, and N is the average number of prints per day for the stocks in a given subgroup.