Order Shredding, Invariance, and Stock Returns

Albert S. Kyle
University of Maryland
akyle@rhsmith.umd.edu

Olga A. Obizhaeva
London School of Economics
O.Obizhaeva@lse.ac.uk

Anna A. Obizhaeva
New Economic School
aobizhaeva@nes.ru

First Draft: September 26, 2013
This Draft: November 17, 2017

Abstract

We introduce a new structural model of stock returns generating process. The model assumes that stock prices change in response to buy and sell bets arriving to the market place as predicted by market microstructure invariance. These bets are shredded by traders into sequences of transactions according to some bet-shredding algorithms. Arbitrageurs take advantage of any noticeable returns predictability, and market makers clear the market. This structural model is calibrated to match empirical time-series and cross-sectional patterns of higher moments of returns. We find that historical idiosyncratic kurtoses of inactively traded stocks are usually higher than that of actively traded stocks, whereas idiosyncratic skewness is positive and stable across stocks, but decrease over time. We calibrate implied hard-to-observe parameters of bet-shredding algorithms using the method of simulated moments, analyse its properties, and show how much shredding has increased over time.

Keywords: market microstructure, invariance, kurtosis, skewness, rare events, fat tails, order shredding, order anticipation, arbitrage.

Introduction

There is a long-standing debate on what is a good way to model security price dynamics. It is crucial for our understanding of financial markets. Progress has been made in this important area, but there is still no fully satisfactory answer as to the mechanism of how returns process is generated. We propose a novel structural model for price dynamics within the paradigm of market microstructure invariance, developed recently by Kyle and Obizhaeva (2016*c*) and found to be successful in explaining a number of empirical regularities in the data.

It is known that empirical price processes depart from the Brownian motion, and price changes are not distributed as normal random variables. Several alternative models have been proposed in the literature. Mandelbrot (1963) suggests that price changes may be better described by a stable Pareto distribution with fat tails. Mandelbrot and Taylor (1967) and Clark (1973) propose that price processes seem to be closely related to the Brownian motions that evolves not in calendar time but rather in some business time, which is linked to either arrival of transactions or trading volume, respectively. Jones, Kaul and Lipson (1994), Hasbrouck (1999), Ané and Geman (2000), Andersen et al. (2015) study what business clock best fits the data. In comparison to these approaches, our structural model of returns dynamics comes from explicit modelling of how traders trade in real financial markets.

The backbone of our model is the arrival process of investment ideas, or bets, placed by fundamental traders into the market. This process has been earlier calibrated within the microstructure invariance paradigm by Kyle and Obizhaeva (2016*c*) who suggested that bets arrive according to a stochastic process with an expected arrival rate per day approximately proportional to the 2/3 exponent of trading volume and volatility, and the distribution of bet sizes closely resemble log-normal random variables with log-variance of 2.53. This large log-variance implies frequent arrivals of very large bets. We assume that traders execute bets by splitting them into sequences of transactions according to some bet-shredding algorithm in order to reduce price impact; we model price impact in response to each transaction as suggested by invariance-based market impact model. We also introduce arbitrageurs who implement order anticipation algorithms based on predictive models to detect execution of large bets and trade ahead of them with hope to make some money. Market makers clear the market.

The core idea of market microstructure invariance is that business time runs faster in liquid markets and slower in illiquid markets, whereas a trading game itself that traders play remains invariant. Our structural model ultimately differs across stocks and time periods, because it is based on different arrival processes of bets. We also calibrate bet shredding parameters using the method of simulated moments in order to match the cross-sectional and time-series variation in empirical moments of stock returns.

We update the evidence on cross-sectional and time-series properties of moments of daily U.S. stock returns using the Center for Research in Security Prices (CRSP) database. We find that idiosyncratic excess kurtoses tend to be positive and decrease with trading activity of stocks; the ratio of idiosyncratic kurtosis for the median least active stocks to that of the most active stocks is almost always greater than one, but this difference becomes less pronounced over time. The total kurtosis without any adjustment for market returns is also larger for the less active stocks; these patterns reverse over during market crashes, when kurtosis of liquid stocks becomes bigger relative to kurtosis of illiquid stocks, possibly due to staleness of prices. The idiosyncratic skewness does not exhibit any distinctive cross-sectional patterns and fluctuates over time around a small positive value, often dropping to negative values during market crashes.

Our calibration allows us to discuss the properties of implied bet shredding parameters. Under the assumption that traders target a fixed proportion of overall expected trading volume, we find that traders target a bigger proportion when executing bets in less liquid securities. We also find that bet-shredding has intensified over time, and now traders choose to execute bets over two or three times longer horizons than in 1950s. The prevalence of shredding in modern markets have been also documented empirically in Kyle, Obizhaeva and Tuzun (2016), Angel, Harris and Spatt (2015), and Garvey, Huang and Wu (2017). Bet shredding is also optimal for traders who seek to minimize transaction costs, as shown theoretically by Kyle, Obizhaeva and Wang (2017). Our structural model can be used as a vehicle to gain insight into hard-to-observe parameters of trading.

There are two different approaches to modelling securities returns. The first approach, usually preferred by economists, relies on calibration of structural equilibrium models in order to make sure that models are internally consistent with market clearing and strategic optimizing behavior of traders; the example is a structural framework of Campbell and Kyle (1993) that helps to model permanent and temporary shocks to prices. The second approach, usually preferred by statisticians and econophysicists, relies on agency-based models that simulate actions and interactions of traders to study their effects on the system as a whole, but often assume mechanic—rather than driven by economic incentives—order placement strategies and price formation process; examples include Cont and Bouchaud (2000), Farmer, Patelli and Zovko (2005), Cont, Stoikov and Talreja (2010), and Ladley (2012), among others.

Our model is a combination of these two approaches, taking the best of both of them. On the one hand, we pay careful attention to the modelling of how people trade as in agency-based models. Indeed, our groups of market participants closely resemble the classification of Kirilenko et al. (2017), who provide a micro-level empirical description of the structure of trading in the market of the E-mini S&P 500 futures during the Flash crash on May 6, 2010. On the other hand, each part of our model is guided by the insights of existing theories of financial economics. Bets arrive according to general invariance predictions, which one can derive within a number of equilibrium models such as the dynamics model of Kyle and Obizhaeva (2017) and the one-period model of Kyle, Obizhaeva and Wang (2017). Bet-shredding algorithms are similar to optimal trading strategies suggested by the literature on optimal execution, such as Bertsimas and Lo (1998), Almgren and Chriss (2000), and Obizhaeva and Wang (2013) among others. Arbitrageurs insure that prices follow a martingale and markets are efficient. Market makers insure that markets clear.

This paper is organized as follows. Section I presents empirical analysis of time-series and cross-sectional properties of moments for returns of the U.S. stock market. Section II describes a structural model of returns dynamics based on market microstructure invariance with bet shredding and arbitrage trading. Section III discuss its calibration and properties of implied parameters. Section IV concludes.

1 Moments of Daily Returns: Empirical Analysis

1.1 Data

We examine cross-sectional and time-series properties of moments of daily U.S. stock returns using the Center for Research in Security Prices (CRSP) database. Only common stocks (CRSP share codes of 10 and 11) listed on the New York Stock Exchange (NYSE), the American Stock Exchange (Amex), NASDAQ, and NYSE Arca in the period of January 1926 through December 2016 are included in the sample. ADRs, REITS, and closed-end funds are excluded.

Estimates of higher moments are very sensitive to large price changes, outliers, and errors in the data. We do not windsorize of the data, because we do not want to eliminate most important large but rare observations. Instead, we carefully clean the data by trying to filter out outliers and errors, while keeping large observations caused by execution of large bets, market crashes, or other events.

First, we adjust for stale prices. For each security, the CRSP mixes two time series. For days with transactions, the database reports last transaction prices at the close. For days with no transactions, the database reports averages of bid and ask prices, marking these averages with a negative sign; these observations are often not representative of true prices, at which traders could actually transact during that day. The mixture two price series often leads to large temporary deviations in the composite series. For example, for the six days from May 17, 2010 to May 24, 2010, one finds the following prices in the CRSP for the stock of the firm Ikonics: \$7.1, \$6.52, -\$7.225, -\$12.76, -\$7.07, and \$6.81; the three negative prices mean that there were no transactions at these three days and the average bid-ask prices are reported instead of actual transaction prices. If one would simply change their negative signs into positives sign and calculate time-series of returns, then he will get -8, 11, 77, -45, and -4 percents with large positive price change followed by large negative price changes in the middle of the sample. At the same time, Yahoo Finance reports $7.1, $6.52, $6.52, $6.52, $6.52, and $6.81 for the same days implying returns of -8, 0, 0, 0, and 4 percents. The two time series will have very different estimates of moments, especially for higher moments such as kurtosis. To circumvent this problem, we use only transaction prices when available, accumulate returns from the very last transaction price reported, and assign returns of zero to all days with no transactions.

Second, there remain many large zigzag price changes in the sample. It is usually unclear whether these are actual prices that we need to keep or errors that we need to eliminate. As describe in Fischer (1963), the process of creating the CRSP database required a lot of effort and involved a lot of data cleaning. Some errors though may still exist due to mistakes in original data collected by exchanges, incorrect conversion of the data from paper books into electronic databases, inconsistent adjustment for splits and dividends, confusion with tickers, inaccurate treatment of trades in error accounts that are often cancelled within a few days, and many other reasons. We checked manually whether large zigzag price deviations in the CRSP coincide with price patterns in other datasets or whether they can be attributed to some events. Since unexplained temporary price swings occur especially often in the earlier pre-war part of the sample, we choose to focus on the data from January 1950 to December 2016.

Third, we eliminate daily observations with fewer than 100 shares traded, because transaction prices on these days may also be not representative of true prices. Small trades may be used as vehicles for side payments between traders, soft commissions, or transactions by market makers who are required to maintain some minimal trading activity in illiquid stocks.

Finally, we exclude stocks with more than fifteen no-trade days in a month and daily volatility of less than one percent. We also exclude stocks with the median of prices being less than $5, because estimates of their returns moments are very unstable, as errors are especially critical for these stocks.

We excluded about 45% of observations from the original sample. The final sample includes 1,576,834 observations for 1,089 months and 19,922 stocks. The number of stocks vary significantly throughout the sample. Initially, there were only NYSE stocks. The number of stocks rose steadily from 500 stocks in 1926 to 1,100 by 1962, then jumped to about 2000 in July 1962 and 5000 in November 1982, when the Amex and NASDAQ stocks were included into the sample, respectively. The number of stocks slightly declined after the market crash of October 1987 and increased during the dot-com bubble 1995 though 2000, peaking at 7300 in 1997. Afterwards, the number of stocks dropped, and it is equal to about 4000 at the end of the sample.

1.2 Estimation of Moments

The estimate moments of log-returns are known to be sensitive to outliers. We next obtain these estimates using robust estimation methods.

We modify the sample estimates of higher moments that usually use the sample estimates of means and standard deviations and that are prone to several biases. First, the sample means introduce forward-looking biases by making returns look less volatile than they are in reality. In our estimation of higher moments, we instead assume that daily stock returns have zero mean.

Second, the sample standard deviation tends to be overestimated during volatile periods, and these biased estimates of volatility in turn make the sample estimates of kurtosis underestimated. We assume means of zero instead sample means and pre-estimate volatility over the previous three-month period, using one of the robust iterative estimation methods; we also consider only three-month periods with more than fifteen non-zero observations of returns and average price above $5. We first estimate volatility over the entire three-month sample, then exclude observations with absolute values bigger than two sigma, estimate volatility again and repeat this procedure until either the difference in subsequent volatility estimates becomes less than one basis point or the number of excluded outliers exceeds five percent of the original sample. These are conservative measure of volatility robust to outliers. For robustness, we also consider volatility estimated using Inter Quantile Range methods (IQR-α methods), as proposed by Aucremanne(2004) and Kimber(1990), respectively, as well as Median Absolute Deviation methods (MAD-β methods), as proposed by Iglewicz and Hoaglin (1993) and Hampel(1974); all results (not reported) are qualitatively and quantitatively similar to our main reported findings.[1]

For each month and each stock, we then calculate the estimates of skewness and kurtosis using the formulas for sample moments but replacing sample means and sample standard deviations with our robust estimates. We apply this procedure for both the sample of returns and the sample of idiosyncratic returns, obtained by subtracting the contemporaneous values of index returns under the assumption that all stocks' betas are equal to one.

1.3 Time-Series and Cross-Section of Empirical Moments

To examine empirically cross-sectional patterns, we split all stocks in ten groups based on daily trading activity, an important characteristic of securities reflecting the speed with which markets operate and levels of liquidity. Trading activity is defined as the product of dollar volume

[1] In IQR-α method, volatility is estimated on reduced sample $[P_{25} - \alpha \cdot [P_{75} - P_{50}], P_{75} + \alpha \cdot [P_{50} - P_{25}]]$, where P_x denote the percentile x, with most outliers excluded ($\alpha = 3$ and $\alpha = 1.5$). In MAD-β method, volatility is estimated on entire sample excluding observations with $|M_i| > \beta$, where $M_i = 0.6745 \cdot (x_i - med(X))/MAD$ and $MAD = med(|x_i - med(X)|)$ ($\beta = 3$ and $\beta = 2$).

and volatility and represent the total amount of risk transferred per day. For each stock and each month, we calculate trading activity as the product of the average daily dollar volume and volatility over the previous three months. We then sort all stocks each month into ten groups based on trading activity. The breakpoints are chosen to be 30th, 50th, 60th, 70th, 75th, 80th, 85th, 90th, 95th of the NYSE traded stocks. Group 1 consists of least actively traded stocks. Group 10 consists of most actively traded stocks.

Table 1 presents a detailed time-series and cross-sectional summary statistics for high moments of idiosyncratic daily returns. The medians of sample moments (volatility, skewness, and kurtosis) are shown for seven decades between 1950 and 2016 and for trading activity groups 1, 3, 5, 8, and 10.

Figure 1 shows the monthly time series of the 12-month moving averages of the sample medians of sample kurtosis of idiosyncratic daily stock returns for the same trading activity groups. The estimates are averaged over a twelve month period to smooth out unstable estimates. Figure 3 shows similar moving averages of monthly kurtosis estimates for daily stock returns without any adjustment for market movements. We can draw several conclusions from table 1 and figure 1.

First, idiosyncratic kurtoses tend to decrease with trading activity. The daily kurtoses of the least liquid stocks are stable, ranging between 6.60 and 8.47 across decades and thus implying fat tails. The daily kurtoses of the most liquid stocks slightly increase over time from 2.66 in decade 1950-1960 to 4.23 in 2010-2016; their values remain being close to 3, suggesting that distributions of their daily idiosyncratic returns closely resemble the log-normal. Figure 2 reveals similar patterns. The figure shows that, depicted by the solid horizontal line, the ratio of idiosyncratic kurtoses of stocks in group 1 to kurtoses of stocks in group 10 is bigger than one for each month throughout the sample, except for the month of September 2008 when uncertainty reached its peak during the financial crisis. The difference in kurtoses of least and most active stocks becomes less pronounced over time. Similar patterns are observed for kurtoses of total daily returns in figure 4. Ratios of kurtoses of stocks in group 1 to kurtoses of stocks in group 10 drop below one only during a few episodes in 1956, 1962-1963, 1987-1988, and 1993-1994; these breaks might be attributed to Kennedy slide in 1962, market crash in October 1987, and mini-crash in October 1989, respectively.

Second, the monthly time series of the 12-month moving averages of the sample kurtosis medians in figures 1 and 3 are relatively stable over time, but exhibit several significant spikes in May 1962, October 1987, August 1998, September 2008, and August 2011. Even though the events that triggered large price changes are relatively short lived, these spikes continue for twelve months due to our calculations of moving averages using the twelve lags. These spikes correspond to volatile times mentioned above as well as to the LTCM collapse in 1998. During these periods, the idiosyncratic kurtoses continue to be larger for less liquid stocks, but the patterns for kurtosis sometimes flip, and kurtosis of liquid stocks becomes bigger relative to kurtosis of illiquid stocks, possibly due to staleness of their price.

Figure 5 shows the time series of idiosyncratic skewness for the trading activity groups. Idiosyncratic skewness is usually slightly positive, fluctuating between 0.06 to 0.36 across decades and decreasing over time, on average, from 0.26 in decade 1950-1960 to 0.10 in 2010-2016, as shown in table 1. During market dislocations, skewness tends to drop. Skewness does not exhibit any distinctive cross-sectional patterns. It remains to be close to zero, thus suggesting that the distribution of returns is close to a log-normal.

Figure 6 shows monthly time series of 12-month moving average of median sample volatility of idiosyncratic daily for the five trading activity groups with two pronounced spikes during the dot-com bubble in 2000-2001 and financial crisis of 2008-2009.

In what follows, we will propose a structural model of price dynamics and calibrate it to match the cross-sectional and time-series patterns of higher moments in table 1.

2 Invariance-Implied Structural Model of Price Dynamics

In this section, we describe a structural model of stock returns dynamics in financial markets. There are three market participants: traders, intermediaries, and arbitragers. Traders are institutional asset managers and retail investors who arrive to the market with some trading ideas, or bets, and execute these bets by shredding them over time based on bet shredding algorithms. We assume that these bets are generated according to the implications of market microstructure invariance. Intermediaries such as traditional market makers and high-frequency traders clear the market by taking the other side of these transactions. Meanwhile arbitrageurs try to detect large bets of traders in the order flow and profit by trading ahead of them.

2.1 Bets of Traders

We start by describing trading strategies of institutional asset managers and retail investors. These traders submit bets based on either some investment ideas or their needs to rebalance portfolios. Bets move prices and induce volatility. Small bets lead to small price changes, large bets trigger large price changes. Invariance implies a specific structure of order flow, i.e. the number of bets and distribution of their size for different markets.

Consider a stock i at day t with returns volatility σ_{it}, share volume V_{it}, dollar price P_{it}, and trading activity

$$W_{jt} = \sigma_{jt} \cdot P_{jt} \cdot V_{jt}. \tag{1}$$

Let γ_{it} denote the number of bets placed at day t in the market of stock i. Suppose that a sequence of bets executed at day t is Q_{it1}, Q_{it2},..$Q_{it\gamma_{it}}$; each kth bet Q_{itk} is measured in shares, bets are positive for buys and negative for sells, both arriving with equal probabilities of 1/2. Let $\tilde{Q}_{it}$ denote a random variable whose probability distribution represents the signed size of bets and let $\tilde{\gamma}_{it}$ denote a random variable whose probability distribution represents the expected arrival rate.

Kyle and Obizhaeva (2016*b*) calibrate these distributions using the sample of portfolio transitions executed over the period 2001 through 2005 in the U.S. stock market as the main benchmark sample. As the first-order approximation, they find that $|\tilde{Q}_{it}|$ is well described by a log-normal distribution with log-variance $\sigma_Q^2 = 2.53$ and $\tilde{\gamma}_{it}$ is a Poisson variable with the mean $\bar{\gamma}_{it}$; the means of both of these random variables vary across days t and stocks i,

$$\bar{\gamma}_{it} = 85 \cdot \Big[\frac{W_{it}}{(0.02)(40)(10^6)}\Big]^{2/3}. \tag{2}$$

$$\ln\Big[\frac{|\tilde{Q}_{it}|}{V_{jt}}\Big] \approx -5.71 - \frac{2}{3} \cdot \ln\Big[\frac{W_{it}}{(0.02)(40)(10^6)}\Big] + \sqrt{2.53} \cdot \tilde{Z}, \qquad \tilde{Z} \sim N(0,1). \tag{3}$$

The 2/3 exponents in these formulas are implications of invariance; the constants 85, −5.71, and 2.53 are calibrated from the data. For the benchmark stock with daily volatility $\sigma = 0.02$, volume $V = 10^6$, and price $P = 40$, for example, there are on average 85 bets per day, their median dollar size is $exp(-5.71) \cdot V \cdot P$ or \$132,000, and their average dollar size is $exp(-5.71 + 0.5\sigma_Q^2) \cdot V \cdot P$ or \$470,000. Both the number of bets $\tilde{\gamma}_{it}$ and their size $|\tilde{Q}_{it}|$ increase with dollar volume and returns volatility.

Intermediaries take the other side of these bets by setting market clearing prices.[2] Kyle and Obizhaeva (2016*b*) analyse by how much each bet on average moves prices and calibrate several price impact models. The first model is the linear price impact model. According to its log-linear version, buying or selling Q shares of a stock with a current stock price P moves the price on average by $\Delta P(Q)$ such that

$$\ln\left(1 + \frac{\Delta P(Q)}{P}\right) = \frac{\sigma_{it}}{0.02}\left(\frac{\bar{\kappa}}{10^4} \cdot \left[\frac{W_{it}}{(0.02)(40)(10^6)}\right]^{-1/3} + 2 \cdot \frac{\bar{\lambda}}{10^4} \cdot \left[\frac{W_{it}}{(0.02)(40)(10^6)}\right]^{1/3} \frac{Q}{(0.01)V_{it}}\right), \tag{4}$$

where $\bar{\kappa} = 8.21$ and $\bar{\lambda} = 2.50$ are calibrated from the data and exponents −1/3 and 1/3 are implications of invariance. The first model is the square root price impact model. According to its log-linear version, buying or selling Q shares of a stock with a current stock price P moves the

[2]Under the assumption that the volume multiplier $\zeta = 2$, as consistent with our assumption that intermediaries take the other side of these bets, and the portfolio transition size multiplier $\delta = 1$.

price on average by $\Delta P(Q)$ such that

$$\ln\left(1+\frac{\Delta P(Q)}{P}\right)=\frac{\sigma_{it}}{0.02}\left(\frac{\bar{\kappa}}{10^4}\cdot\left[\frac{W_{it}}{(0.02)(40)(10^6)}\right]^{-1/3}+2\cdot\frac{\bar{\lambda}}{10^4}\cdot\left[\frac{Q}{(0.01)V_{it}}\right]^{1/2}\right), \tag{5}$$

where $\bar{\kappa}=2.08$ and $\bar{\lambda}=12.08$ are calibrated from the data and exponents $-1/3$ and $1/2$ are implications of invariance.

Equations (2) and (3) describe the order-flow process for traders. Equations (4) and (5) describe how intermediaries update prices in response to each bet. Combining price impact of all bets executed during the day, one can calculate implied daily price changes. The set of these equations thus describe a basic structural model for daily returns, as implied by invariance.

2.2 Price Changes Upon Execution of One Bet

We next examine moments of price changes induced by one bet. Since buys and sells arrive with equal probabilities, the distribution of signed bet sizes $\tilde{Q}_{it}$ is symmetric, and all of its odd moments are equal to zero. For example, $E\left[\tilde{Q}_{it}\right]=0$ and $E\left[\tilde{Q}_{it}^3\right]=0$.

Since the distribution (3) of unsigned bet size $|\tilde{Q}_{it}|=\exp(\mu_Q+\sigma_Q\cdot\tilde{Z})$ is a log-normal with a log-mean of μ_Q and a log-variance of $\sigma_Q^2=2.53$, its moments can be calculated as,

$$E\left[|\tilde{Q}_{it}|^p\right]=\int q^p\cdot\frac{1}{q}\cdot\frac{1}{\sqrt{2\pi\sigma_Q^2}}\exp\left(-\frac{(\ln(q)-\mu_Q)^2}{2\sigma_Q^2}\right)dq=e^{p^2\sigma_Q^2/2+p\mu_Q}. \tag{6}$$

This implies the kurtosis of price changes upon execution of a bet. For the linear price impact model, it is equal to kurtosis of a bet size itself,

$$\text{kurt}\left[\Delta P(\tilde{Q}_{it})\right]=\text{kurt}\left[|\tilde{Q}_{it}|\right]=\frac{\mathbb{E}\left[|\tilde{Q}_{it}|^4\right]}{\mathbb{E}\left[|\tilde{Q}_{it}|^2\right]^2}=e^{4\sigma_Q^2}=22,000. \tag{7}$$

For the square root price impact model, it is equal to

$$\text{kurt}\left[\Delta P(\tilde{Q}_{it})\right]=\text{kurt}\left[|\tilde{Q}_{it}|^{1/2}\right]=\frac{\mathbb{E}\left[|\tilde{Q}_{it}|^2\right]}{\mathbb{E}\left[|\tilde{Q}_{it}|\right]^2}=e^{\sigma_Q^2}=12. \tag{8}$$

These values are much larger than kurtosis of 3 for a normal distribution, especially for the linear model.

2.3 Price Changes upon Execution of Bet Sequences with No Bet Shredding

Daily price change ΔP is equal to the sum of all price changes in response to execution of independent and identically distributed bets. If there are γ bets executed in day t and stock i, then kurtosis of daily returns is

$$\operatorname{kurt}\left[\Delta P|\tilde{\gamma}_{it}=\gamma\right]=\operatorname{kurt}\left[\sum_{k=1}^{\gamma}\Delta P(Q_{kit})\right]=\frac{\operatorname{kurt}\left[\Delta P(\tilde{Q}_{it})\right]}{\gamma}, \tag{9}$$

where $\operatorname{kurt}\left[\Delta P(\tilde{Q}_{it})\right]$ is defined in equations (7) and (8).

To find unconditional kurtosis, we should integrate out γ in equation (9), because the number of bets executed per day is a random variable. If no bet arrives, then we should not update our estimates of kurtosis. The kurtosis of the random sum of random variables with expected Poisson arrival rate $\bar{\gamma}_{it}$ is given by

$$\operatorname{kurt}[\Delta P]=E_{\gamma}\left(\operatorname{kurt}\left[\Delta P|\tilde{\gamma}_{it}=\gamma\right]\right)=\sum_{j=1}^{+\infty}\left[\frac{\operatorname{kurt}\left[\Delta P(\tilde{Q}_{it})\right]}{j}\cdot\frac{\bar{\gamma}_{it}^{j}}{j!}\cdot e^{-\bar{\gamma}_{it}}\right]. \tag{10}$$

The infinite sum $\sum_{j=1}^{+\infty}\left[\gamma^{j}/j!\right]$ is a converging series, a series $\{1/j\}$ is a bounded from above, monotone sequence, and $\operatorname{kurt}\left[\Delta P(\tilde{Q}_{it})\right]$ is a constant. Applying Abel's convergence test, we find that the infinite sum (10) converges, though it does not have a close form solution.

It is possible to derive the lower bound for the unconditional kurtosis using Jensen's inequality. Indeed,

$$\operatorname{kurt}[\Delta P]=E_{\gamma}\left(\frac{\operatorname{kurt}\left[\Delta P(\tilde{Q}_{it})\right]}{\tilde{\gamma}_{it}}\right)\geq\frac{\operatorname{kurt}\left[\Delta P(\tilde{Q}_{it})\right]}{E_{\gamma}\left(\tilde{\gamma}_{it}\right)}=\frac{\operatorname{kurt}\left[\Delta P(\tilde{Q}_{it})\right]}{\bar{\gamma}_{it}}. \tag{11}$$

The lower bound is equal to the kurtosis of daily returns (9) conditional of the assumption that the number of bets $\tilde{\gamma}_{it}$ coincides with the average arrival rate $\bar{\gamma}_{it}$.

Our simulation analysis shows that this lower bound provides a good approximation for the daily kurtosis of most stocks, as implied by a structural model. The differences created by uncertainty in the Poisson arrival rates and non-linearities of log-returns are insignificant. For a stock with median dollar volume and returns variance in each of the ten trading activity groups, we run 1000 Monte-Carlo simulations and calculate the average theoretical kurtosis with its standard errors. The simulations are done based on a structural model of price process with bet arrival rate in equation (2), distribution of bet sizes in equation (3), and price impact model (4). We also calculate the lower bound using equation (11).

Table 2 shows that the lower bound tracks closely the average kurtosis for all groups, except the group of least actively traded stocks. The percentage differences in the series of two estimates are 29%, 3%, 2%, 1%, and 0% for groups 1, 3, 5, 8, and 10, respectively. A large difference for the first group may reflect an upward bias in theoretical estimates of kurtosis. Since the arrival rate for this group is only 4 bets per day, some simulated days have no bets and their exclusion from calculations of the average introduces a bias. The bet arrival rates for other groups range from 23 to 232, and this effect is less pronounced. As long as the arrival rate of bets is not too low, the lower bound is a reasonable proxy for kurtosis of daily returns. We get the following approximation,

$$\text{kurt}\,[\Delta P] \approx \frac{\text{kurt}\left[\Delta P(\tilde{Q}_{it})\right]}{\bar{\gamma}_{it}}. \tag{12}$$

Using equations (7) and (8), the lower bound for daily kurtosis is equal to $22,000/\bar{\gamma}_{it}$ and $12/\bar{\gamma}_{it}$ for the linear and square root models, respectively.

Kurtosis of price changes per each bet is the same across stocks, but the number of bets per day is larger for more liquid stocks. Therefore daily returns of more liquid stocks have lower kurtosis. The number of bets $\bar{\gamma}_{it}$ per day increases with trading activity at a rate of 2/3 in equation (2). Equation (12) then implies that kurtosis decreases with the trading activity approximately at the same rate, i.e., with 2/3 power of the trading activity. In table 2, for example, the ratio of kurtosis of most inactively traded stock to kurtosis of most actively traded stock is about 77 (= 7214/95); it is similar to the ratio of their trading activities in 2/3 power equal to 59 (= $(3600/8)^{2/3}$).

Similar intuition suggests that kurtosis decrease with tenor of returns for a given security. For example, kurtosis of weekly returns is expected to be lower than kurtosis of daily returns, which in turn is expected to be lower than kurtosis of one-minute returns.

Our basic structural model implies the values of kurtosis that are too high relative to empirical estimates. The average theoretical kurtosis in table 2 ranges between 95 and 7,214, whereas empirical estimates in table 1 do not exceed 8.47, the level of average kurtoses for stocks in group 1 for decade 1960-1970.

2.4 Price Changes upon Execution of Bet Sequences with Bet Shredding

So far we have assumed each bet is executed instantaneously. In reality, traders shred orders and execute them over time in sequences of transactions to reduce transaction costs. Bet shredding smooths out spikes in price dynamics and tends to make returns kurtoses smaller. We next consider several modifications of our basic model that incorporate order shredding and arbitrage trading. These models are more realistic and more flexible in their ability to match empirical estimates.

Traders decide on “target” inventories and bets based on either their private information or inventories shocks. Then, they gradually adjust actual inventories towards their targets. Let S^*_{it} denote the cumulative target order imbalances for stock i at the end of day t, calculated as the signed sum of all *bets* placed into the market place by that time,

$$S^*_{it} = \sum_{m \le t} Q_{imk}. \tag{13}$$

Suppose next that each bet Q_{imk} is shredded into a sequence of transactions $x_{imk}(s)$, where s is a day count in execution package.

Let S_{it} denote cumulative realized order imbalance for stock i at the end of day t, calculated as the signed sum of all *transactions* placed into the market by that time.

$$S_{it} = \sum_{m \le t, s \le t} x_{imk}(s). \tag{14}$$

The structural model of trading (13) and (14) is consistent with the equilibrium strategies in a continuous-time model of smooth trading of Kyle, Obizhaeva, Wang (2016). In that model, symmetric, relatively overconfident, oligopolistic informed traders calculate target inventories based on how their own estimates of the long-term dividend growth rate differ from the estimates of other traders. Since the market offers no instantaneous liquidity for block trades, each trader only partially adjusts his inventory in the direction of a target inventory; the rate of adjustment is determined by the deep parameters of the model, it is larger when private information decays faster and when there is more disagreement between traders.

The difference between the time series of S^*_{it} and S_{it} depends on the specifics of bet shredding algorithms. Bet shredding algorithms are not directly observable. We assume that each algorithm is characterized by two main decisions. For each bet, traders first choose execution horizon and then parameters of shredding method. We consider several alternative specifications.

First, traders determine an appropriate execution horizon T_{itk} for each bet Q_{itk}. For example, traders may target a fixed time horizon t, say one day,

$$\text{Method-}T(t): \qquad T_{itk} = t. \tag{15}$$

We refer to this algorithm as "Method-$T(t)$"; for example, "Method-$T(1)$" or "Method-$T(5)$" correspond to cases when all trades are executed their bets over one day or one week.

Traders may also target a small fraction η, say equal to 5%, of expected contemporaneous volume $T_{itk} \cdot V_{it}$ or $T_{itk} \cdot \bar{\gamma}_{it} \cdot E\left[|\tilde{Q}_{it}|\right]$,

$$|Q_{itk}| = \eta \cdot T_{itk} \cdot \bar{\gamma}_{it} \cdot E\left[|\tilde{Q}_{it}|\right]. \tag{16}$$

This implies the execution horizon that is linearly proportional to bet size,

$$\text{Method-}V(\eta): \qquad T_{kit} = \frac{|Q_{kit}|}{\eta \cdot \bar{\gamma}_{it} \cdot E\left[|\tilde{Q}_{it}|\right]}. \tag{17}$$

We refer to this algorithm as "Method-$V(\eta)$"; for example, "Method-$V(0.05)$" or "Method-$V(0.10)$" for execution algorithms targeting 5 percent and 10 percent of daily volume, respectively.

Traders may also target to induce a small fraction η, say equal to 5%, of expected returns variance $T_{kit} \cdot \sigma^2_{it}$ under the assumption that each transaction is expected to move price by $\lambda \cdot |Q_{kit}|$,

$$(\lambda \cdot |Q_{kit}|)^2 = \eta \cdot T_{kit} \cdot \gamma_{it} \cdot \lambda^2 \cdot E\left[|\tilde{Q}_{it}|^2\right]. \tag{18}$$

This implies that the execution horizon is proportional to the square of bet size,

$$\text{Method-}\sigma^2(\eta): \qquad T_{kit} = \frac{|Q_{kit}|^2}{\eta \cdot \gamma_{it} \cdot E\left[|\tilde{Q}_{it}|^2\right]}. \tag{19}$$

We refer to this algorithm as "Method-$\sigma^2(\eta)$"; for example, "Method-$\sigma^2(0.05)$" or "Method-$\sigma^2(0.10)$" for execution algorithms targeting 5 percent and 10 percent of daily volatility, respectively.

In all cases, larger bets are executed over longer period of time. In the third case (17) larger bets are spread over longer periods of time than in the second case (19) and returns distribution is expected to exhibit smaller kurtosis. For the square root impact model, targeting a given fraction of returns variance is equivalent to targeting a given fraction of volume, so we do not consider this case separately.

Next, traders have to choose an appropriate shredding method. We consider two bet shredding methods. Each bet Q_{kit} can be shredded at a uniform rate and executed in equally-sized transactions $x_{kit}(s)$,

$$x_{itk}(s) = \frac{|Q_{itk}|}{T_{itk}}, \qquad s = 1,..T_{itk}. \tag{20}$$

Bertsimas and Lo (2001) find that this simple execution is optimal when a risk-neutral trader needs to execute an order.

Alternatively, each bet Q_{itk} can be shredded at a monotonically decreasing rate,

$$x_{itk}(s) = |Q_{itk}| \cosh(\rho(T_{itk} - s + 0.5)) \frac{2\sinh(0.5\rho)}{\sinh(\rho T_{itk})}, \qquad s = 1,.T_{itk}. \tag{21}$$

where sinh and cosh are the hyperbolic sine and cosine functions. Each day a trader executes some fraction of the remaining part of the bet, determined by parameter ρ. This parameter is related to the speed of information decay, risk aversion, and riskiness of securities. The larger is parameter ρ, the faster the bet is executed. Almgren and Chriss (2000) finds that execution is optimal when a risk-averse trader executes a bet. Similar solution can be also found in Grinold and Kahn (1999).

We choose to focus on simple execution strategies. In reality, execution strategies are more complicated. Execution algorithms are often price dependent, as discussed in Obizhaeva (2012). Other order shredding algorithms are for example discussed in Gatheral and Schied (2013), Schied and Schoeneborn (2009), and Obizhaeva and Wang (2013). If necessary, sophisticated execution strategies may be built into our structural model as well.

The structural model for bet arrival (2) and (3) augmented with specific order shredding algorithm represent the structural model describing the order-flow process. Together with price impact model, they allow to construct implied time-series of prices. In what follows, we consider linear price impact rule.

2.5 Price Dynamics with Shredding and Arbitrageurs

Bet shredding introduces positive autocorrelation in stock return process and makes future price changes predictable. For example, execution of a large buy bet is expected to inject a positive trend into the price dynamics, while execution of a large sell bet induces a downward price dynamics. Arbitrageurs notice that prices are not martingales and construct order anticipation algorithms to detect execution of orders.

We next describe how to model trading by arbitrageurs. If intermediaries observed target bet imbalances, they would set prices according to their market clearing rule,

$$P^*_{it} = \lambda_{it} \cdot S^*_{it}, \tag{22}$$

and price changes would be unpredictable. In reality, intermediaries may at best identify only actual signed order imbalances S_{it} and set prices as,

$$\hat{P}_{it} = \lambda_{it} \cdot S_{it}. \tag{23}$$

To the extent that unexecuted order imbalance $S^*_{it} - S_{it}$ are predictable based on past information, these price process is not a martingale.

Arbitrageurs build a model to forecast $S^*_{it} - S_{it}$ and trade $\mathbb{E}_t\{S^*_{it} - S_{it}\}$ at day t. When target order imbalances are higher than actual order imbalances, arbitrageurs buy ahead of other traders. When target inventories are lower than actual inventories, arbitrageurs sell ahead of other traders. Market makers set clearing prices based on the aggregate order flow of both traders and arbitrageurs,

$$P_{it} = \mathbb{E}_t\{P^*_{it}\} = \lambda_{it} \cdot S_{it} + \lambda_{it} \cdot \mathbb{E}_t\{S^*_{it} - S_{it}\} = \lambda_{it} \cdot \mathbb{E}_t\{S^*_{it}\}. \tag{24}$$

Trading by arbitrageurs restore martingale properties of stock prices and makes price process $P_{it} = \mathbb{E}\{P^*_{it}\}$ a martingale based on arbitrageurs' filtration. Essentially, the price is set based on the market's forecasts of current target imbalances.

Our structural model is flexible to be consistent with various predictive models of arbitrageurs. We suppose that arbitrageurs know daily volatility and daily volume of an asset. They are also familiar with all invariance formulas and bet shredding algorithms that traders use. Arbitrageurs thus can simulate hypothetical bet arrival process and how bets are shredded into sequences of transactions. Then, they can perform a large estimation on the simulated sample to build a model for forecasting unexecuted order imbalances.

This procedure can be summarized as follows,

1. Simulate N paths of bet histories for an asset with volume V_{it} and volatility σ_{it} based on formulas (2) and (3);

2. Using the conjectured parameters of bet shredding algorithm, aggregate bets and transactions, calculating histories of target imbalances and actual imbalances, $S^*_{it,n}$ and $S_{it,n}$ for each of simulated paths $n = 1,..N$;

3. Run a rolling-window predictive regression for unexecuted imbalances with k lags of linear and quadratic terms of realized past imbalances,

$$\mathbb{E}_t\{S^*_{it,n} - S_{it,n}\} = \alpha + \sum_{j=1}^{k} \beta_{1j} \cdot S_{i,t-j,n} + \sum_{j=1}^{k} \beta_{2j} \cdot S^2_{i,t-j,n} + \epsilon_{tn}, \qquad t = 1,..T, n = 1,..N, \tag{25}$$

   to estimate coefficients $\hat{\beta}_{1j}$ and $\hat{\beta}_{2j}$, $j = 1,..k$. For example, we use $k = 5$ as our benchmark model, i.e. an arbitrageur using information on actual inventories over the previous week.

Equipped with estimated model $\hat{\beta}_{1j}$ and $\hat{\beta}_{2j}$, $j = 1,..k$, arbitrageurs construct forecasts based on current information about past order imbalances $S_{i,t-j}$, $j = 1,..k$, as

$$\mathbb{E}_t\{S^*_{it} - S_{it}\} = \alpha + \sum_{j=1}^{k} \hat{\beta}_{1j} \cdot S_{i,t-j} + \hat{\beta}_{2j} \cdot S^2_{i,t-j}. \tag{26}$$

This is a model for forecasting an unexecuted order imbalances.

In what follows, we mostly apply bet shredding method that targets a given fraction of daily volume, split all bets into equally-sized transactions, and assume a linear price impact function.

2.6 Properties of Simulated Returns

We first illustrate our structural model using the example of a hypothetical benchmark stock with price P of \$40 per share, daily volume V of one million shares, and daily volatility σ of 2% per day. This benchmark stock would belong to the bottom tercile of S&P 500.

We simulate 1,000 paths of 90-day bet arrival histories for the benchmark stock using formulas (2) and (3). We then apply several bet shredding algorithms by first selecting the execution horizon depending on the fraction of daily volume targeted and second by shredding each bet into a sequence of equally-sized transactions. The execution of some packages extends beyond the boundary of a 90-day sample. We cut tails of these unfinished packages at the end of each sample path, multiply remaining sequences by 1 or -1 with equal probabilities to model buy and sell orders, and insert them into the beginning of the same sample path. This mimics a typical situations when some of large bets arrived in the past are continuing to get executed at the beginning of selected sample paths.

We then estimate forecasting model (25) of arbitrageurs, who seek to predict unexecuted bet imbalances at each point of time using the last five realized bet imbalances and their squares. This estimation is done on the entire simulated sample on a rolling-window basis.

Table 3 reports the results for the three bet shredding algorithms with $\eta = 1\%$, $\eta = 5\%$, and $\eta = 10\%$. The lower is fraction η of volume targeted in the execution, the more execution is extended over time, the more past imbalances are autocorrelated with current unexecuted imbalances, and the larger are estimated coefficients. For example, when $\eta = 1\%$, the coefficients are 1.98, 0.97, 1.02, 1.25, and 5.32. When $\eta = 10\%$, the coefficients are only 0.17, 0.16, 0.19, 0.30, and 0.89. Using these estimates, we construct predictive model (26) and price paths using equation 4.

Figure 7 shows the averages, medians, and standard error bounds for returns autocorrelation coefficients at different lags, ranging from one day to forty days for the simulated sample under the assumption that there are no arbitrageurs. The four panels show the results for the cases of $\eta = 1\%$, $\eta = 5\%$, $\eta = 10\%$, and the case with no shredding, i.e. $\eta = \infty$. As expected, when there is no shredding, autocorrelations are equal to zero at all lags. In the other panels, autocorrelations are high at fist lags, decaying with time. The lower is the fraction η of bet shredding algorithm and the longer are execution horizons of large bets, the bigger autocorrelations at first lags are and the slower they decay.

Figure 8 show the same statistics but under the assumption that there are arbitrageurs. Most of the autocorrelation coefficients are now close to zero, since arbitrageurs eliminate most of returns predictability. Based on the linear terms and squared terms, their forecasting model works reasonably well, except for reducing autocorrelations at the boundaries of their forecasting window, which is assumed to have the length of five days in our example.

Table 4 presents the autocorrelations and their standard errors. As before, in panel A when the model has no arbitrageurs, many of the coefficients are statistically bigger than zeros, especially when η is small. In panel B when we introduce arbitrageurs, most coefficients become insignificant. For example, when $\eta = 1\%$, the first-order autocorrelation is equal to 0.696 with standard errors of 0.092 with no arbitrageurs and 0.033 with standard errors of 0.119 with arbitrageurs.

Figures 9 and 10 present distributions of the four moments of simulated returns for the cases without and with arbitrageurs, respectively. There are distributions of the four moments in the four columns. Each of the four rows corresponds to different bet-shredding methods with $\eta = 1\%$, $\eta = 5\%$, $\eta = 10\%$ as well as the case with no shredding. Table 5 reports the summary statistics for these distributions. On both figures, the means and the skewness are centered around zero, since the base model is symmetric for buy and sell orders. The volatility is much lower than initially assumed daily volatility of $\sigma = 2\%$ when there are no arbitrageurs, especially when η is low. Intuitively, bet-shredding converts returns volatility into the price drift. Trading by arbitrageurs "restores" martingale properties of prices and brings volatility back to the assumed levels. For example, when $\eta = 1\%$, the daily volatility of simulated returns is equal to 0.005 with no arbitrageurs and 0.022 with arbitrageurs.

3 Properties of Implied Shredding Parameter

The properties of daily returns depend on the assumptions about parameters of the bet-shredding algorithm. We next use the method of simulated moments and calibrate these parameters to match empirical moments of daily returns.

As before, we assume that traders generate bets according to invariance, design execution to target a given fraction η of expected daily volume, and split bets into equally-sized transactions. Meanwhile, arbitrageurs apply the forecasting model described in section 2.5 and market makers clear the market. We generate $N = 1000$ paths of daily returns. The bet-shredding parameter η is then estimated by matching the kurtoses of simulated returns $\mathrm{kurt}(\Delta P|\eta, n)$ to the empirical estimates of kurtoses $\mathrm{kurt}(\Delta P|\mathrm{Data})$,

$$\eta^* = \mathrm{argmin}_\eta \Big(\frac{\sum_{n=1}^N \mathrm{kurt}(\Delta P|\eta, n)}{N} - \mathrm{kurt}(\Delta P|\mathrm{Data}) \Big). \tag{27}$$

The empirical estimates are taken from table 1 for different trading activity groups and time periods.

Table 6 reports the estimates of implied parameter η for median stocks in the five out of ten trading activity groups and for the seven decades from 1950 to 2017. The table also presents information about trading activity used for simulation of daily returns; its values coincide with statistics reported in table 1. There are two patterns.

First, the implied parameter η decreases over time. For the stocks in group 1, parameter η decreased from 8.875 during 1950–1960 to 5.04 during 2010–2017. For the stocks in group 10, parameter η decreased from 3.225 during 1950–1960 to 1.39 during 2010–2017. This implies that, conditional of bet size, bet shredding increased over time. Similarly, **?** document a significant change in trading patterns in the Trades and Quotes (TAQ) dataset, as the decimalization and use of electronic interfaces has recently led to a significant increase in order shredding; the market for block trades seems almost to have disappeared, and most trading is now dominated by transactions of 100 shares, the minimum lot size. The feature of increased bet shredding implied by our structural model suggests that it has reasonable properties.

Second, the implied parameter η decreases with trading activity W. For example, for the time period 1990 through 2000, η is equal to 8.36, 4.22, 3.46, 2.52, and 1.66 for groups 1, 3, 5, 8, and 10, respectively. For the time period 2010 through 2017, η is equal to 5.04, 2.69, 2.32, 1.80, and 1.39, respectively. This implies that, conditional on bet size, execution of bets is spread over longer periods for more actively traded stocks.

Table 7 shows implied execution horizons for the two periods before and after decimalization. Panel A shows results for 1990-2000. Panel B shows results for 2010-2017. We calculate bet sizes using equation (3) and then calculate implied execution horizons using equation (17) and calibrated bet-shredding parameters $\hat{\eta}$ from table 6. For the median stock in group 1, it takes 2.69, 13.22, 64.86, and 318.23 minutes to execute 4-std, 5-std, 6-std, and 7-std bets during 1990–2000, respectively, and 0.79, 3.86, 18.93, and 92.87 minutes for similar bets during 2010–2017. For the median stock in group 10, it takes 0.17, 0.85, 4.16, and 20.39 minutes to execute 4-std, 5-std, 6-std, and 7-std bets during 1990–2000, respectively, and 0.05, 0.26, 1.29, and 6.32 minutes for similar bets during 2010–2017. The speed of execution increased by a factor of 3.

The inspection of estimates in table 7 reveals that differences in bet-shredding parameters are similar to differences in trading activity in $-1/3$ power. For example, the ratio in parameters η_i and η_j for stock i and j are related to the ratio of their trading activities W_i and W_j as approximately,

$$\frac{\eta_i}{\eta_j} \approx \left(\frac{W_i}{W_j}\right)^{-1/3}. \tag{28}$$

We can also extrapolate these estimates to the overall market with daily trading volume of

$292 billion (futures and stocks combined) and daily volatility of 2 percent, as noted in Kyle and Obizhaeva (2016*a*). Using parameters for the median stock in group 10 during 2010–2017 as the benchmark, equation (28) implies that bet-shredding parameter for the entire U.S. market $\eta_{\text{mkt}} \approx 1.39 \cdot (\frac{W_i}{7{,}141{,}896})^{-1/3} \approx 0.20$, i.e. traders are targeting about 20 percent of expected contemporaneous volume when executing bets in the U.S. market. This is broadly consistent with information in Staffs of the CFTC and SEC (2010b) that the large trader whose trading caused the Flash crash on May 6, 2010, has been targeting 9 percents of contemporaneous volume when executing a bet in the E-mini S&P500 futures market.

4 Conclusions

We propose a new structural model of stock returns dynamics, which is inspired by the recently developed ideas of market microstructure invariance. Traders generate investment ideas, or bets, and execute them by shredding large orders over time to minimize transaction costs, arbitrageurs trade to profit on any detectable trends in prices, and market makers clear the market. Bets are assumed to arrive according to the processes calibrated by Kyle and Obizhaeva (2016*b*); parameters of bet-shredding algorithms are chosen to match empirical moments of stock returns.

Our structural model captures realistically the economics of trading. It is the model of stochastic volatility, because arrival of bets and their sizes are stochastic, and large bets lead to bursts in volume, volatility, and intermediation. The model is flexible in terms of modelling trading behavior of arbitrageurs and bet-shredding algorithms, while precise and grounded in theory in terms of using a specific structure of bet flow from traders and intermediaries. It can be calibrated either to fit the data or to infer the implied parameters of trading, for example, such as hard-to-observe bet-shredding parameters.

We focus mostly on the price dynamics, but the framework also generates quantitative predictions about overall trading volume and order flow generated by different groups of traders. As an extension, it is possible to calibrate the model to match cross-sectional and time-series properties of both stock returns and trading volume, or even some empirical findings about trading by different groups of traders, for example, such as defined in Kirilenko et al. (2010).

References

Almgren, Robert, and Neil Chriss. 2000. "Optimal Execution of Portfolio Transactions." *Journal of Risk*, 3(2): 5–39.

Andersen, Torben G., Oleg Bondarenko, Albert S. Kyle, and Anna A. Obizhaeva. 2015. "Intraday Trading Invariance in the E-Mini S&P 500 Futures Market." Working Paper. Available at http://dx.doi.org/10.2139/ssrn.2693810.

Ané, Thierry, and Helyette Geman. 2000. "Order Flow, Transaction Clock, and Normality of Asset Returns." *The Journal of Finance*, 55(5): 2259–2284.

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.

Bertsimas, Dimitris, and Andrew Lo. 1998. "Optimal Control of Execution Costs." *Journal of Financial Markets*, 1(1): 1–50.

Campbell, John Y., and Albert S. Kyle. 1993. "Smart Money, Noise Trading, and Stock Price Behavior." *The Review of Economic Studies*, 60: 1–34.

Clark, Peter. 1973. "A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices." *Econometrica*, 41(1): 135–155.

Cont, Rama, and Jean-Philipe Bouchaud. 2000. "Herd Behavior and Aggregate Fluctuations in Financial Markets." *Macroeconomic Dynamics*, 4: 170–196.

Cont, Rama, Sasha Stoikov, and Rishi Talreja. 2010. "A Stochastic Model for Order Book Dynamics,." *Operations Research*, 58(3): 549–563.

Farmer, Doyne J., Paolo Patelli, and Ilija I. Zovko. 2005. "The Predictive Power of Zero Intelligence in Financial Markets." *Proceedings of the National Academy of Sciences of the USA*, 102(11): 2254–2259.

Fischer, Lawrence. 1963. "Use of electronic computers in the quality control of financial data." *Presentation at the American Statistical Association.*

Garvey, Ryan, Tao Huang, and Fei Wu. 2017. "Why Do Traders Split Orders?" *Financial Review*, 52(2): 233–258.

Gatheral, Jim, and Alex Schied. 2013. "Dynamical models for market impact and algorithms for optimal order execution." *Handbook on Systemic Risk (eds.: J.-P. Fouque and J. Langsam)*, 579–602.

Hasbrouck, Joel. 1999. "Trading Fast and Slow: Security Market Events in Real Time." Working Paper, New York University.

Jones, Charles M., Gautam Kaul, and Marc L. Lipson. 1994. "Transactions, Volume, and Volatility." *The Review of Financial Studies*, 7(4): 631–651.

Kirilenko, Andrei, Albert S. Kyle, Tugkan Tuzun, and Mehrdat Samadi. 2010. "The Flash Crash: High-Frequency Trading in an Electronic Market." Working Paper, University of Maryland.

Kirilenko, Andrei, Albert S. Kyle, Tugkan Tuzun, and Mehrdat Samadi. 2017. "The Flash Crash: High-Frequency Trading in an Electronic Market." *Journal of Finance*, 72(3): 967–998.

Kyle, Albert S., and Anna A. Obizhaeva. 2016*a*. "Large Bets and Stock Market Crashes." Available at SSRN: `https://ssrn.com/abstract=2023776` or `http://dx.doi.org/10.2139/ssrn.2023776`.

Kyle, Albert S., and Anna A. Obizhaeva. 2016*b*. "Market Microstructure Invariance: Empirical Hypotheses." *Econometrica*, 84(4): 1345–1404.

Kyle, Albert S., and Anna A. Obizhaeva. 2017. "Market Microstructure Invariance: A Dynamic Equilibrium Model." Working Paper, University of Maryland and New Economic School.

Kyle, Albert S., and Anna Obizhaeva. 2016*c*. "Market Microstructure Invariance: Empirical Hypotheses." *Econometrica*, 84(4): 1345–1404.

Kyle, Albert S., Anna A. Obizhaeva, and Tugkan Tuzun. 2016. "Microstructure Invariance in U.S. Stock Market Trades." Working Paper, University of Maryland.

Kyle, Albert S., Anna A. Obizhaeva, and Yajun Wang. 2017. "Smooth Trading with Overconfidence and Market Power."

Ladley, Dan. 2012. "Zero Intelligence in Economics and Finance." *Agent-Based Computational Economics*, 27(2): 273–286.

Mandelbrot, Benoît. 1963. "The Variation of Certain Speculative Prices." *The Journal of Business*, 36(4): 394–419.

Mandelbrot, Benoît, and Howard M. Taylor. 1967. "On the Distribution of Stock Price Differences." *Operations Research*, 15(6): 1057–1062.

Obizhaeva, Anna A. 2012. "Liquidity Estimates and Selection Bias." Working Paper. Available at `https://ssrn.com/abstract=1178722` or `http://dx.doi.org/10.2139/ssrn.1178722`.

Obizhaeva, Anna A., and Jiang Wang. 2013. "Optimal Trading Strategy and Supply/Demand Dynamics." *Journal of Financial Markets*, 16(1): 1–32.

Schied, Alex, and Torsten Schoeneborn. 2009. "Risk aversion and the dynamics of optimal trading strategies in illiquid markets." *Finance and Stochastics*, 13(2): 181–204.

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

Figure 1: Idiosyncratic Kurtosis of Daily Stock Returns for 1950 through 2016.

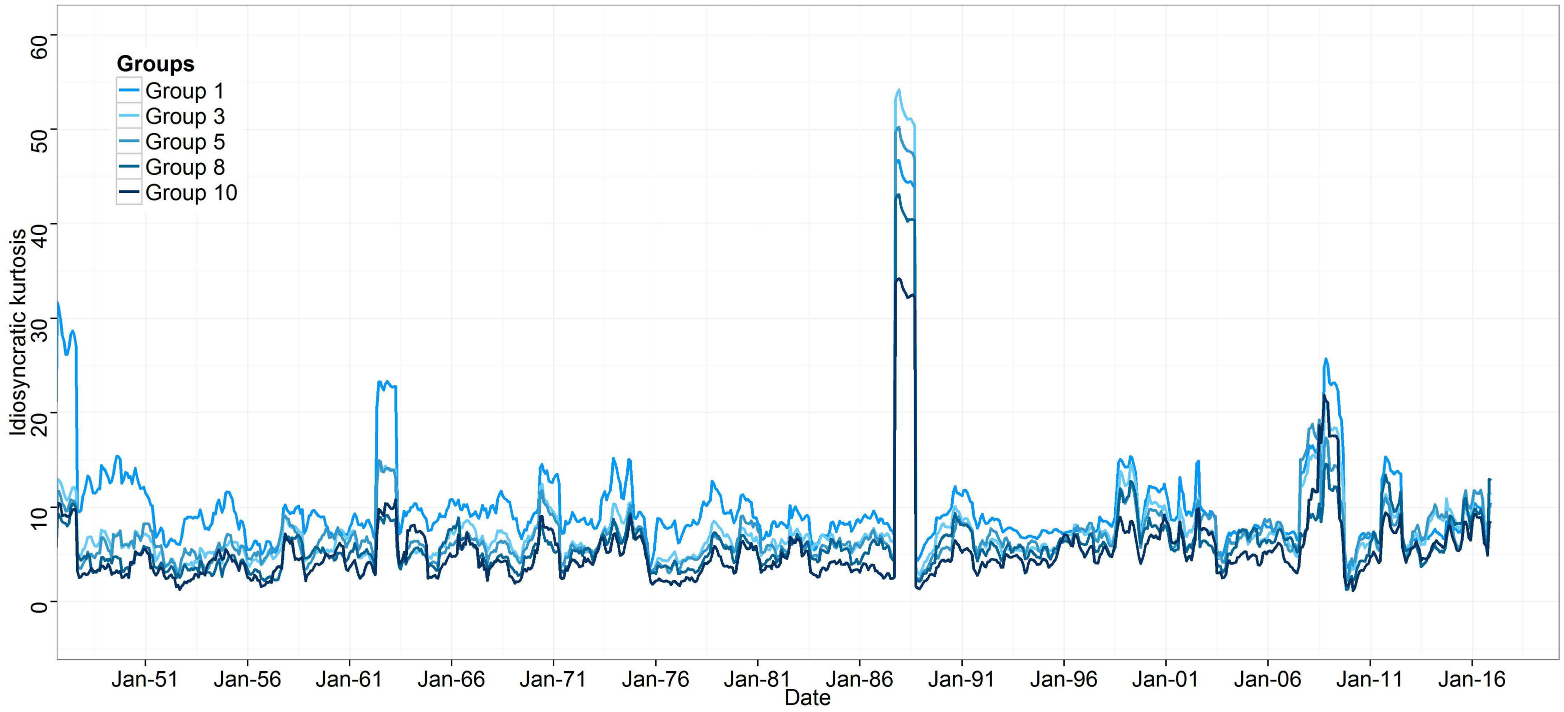

Figure shows five monthly time series of 12-month moving averages of median sample kurtosis for idiosyncratic daily stock returns for each of the five trading activity groups (groups 1, 3, 5, 8, and 10 out of ten groups). Group 1 (10) contains the least (most) actively traded stocks. The period ranges from January 1950 to December 2016.

Figure 2: Ratio of Idiosyncratic Kurtosis (Group 1 to Group 10) for 1950 through 2016.

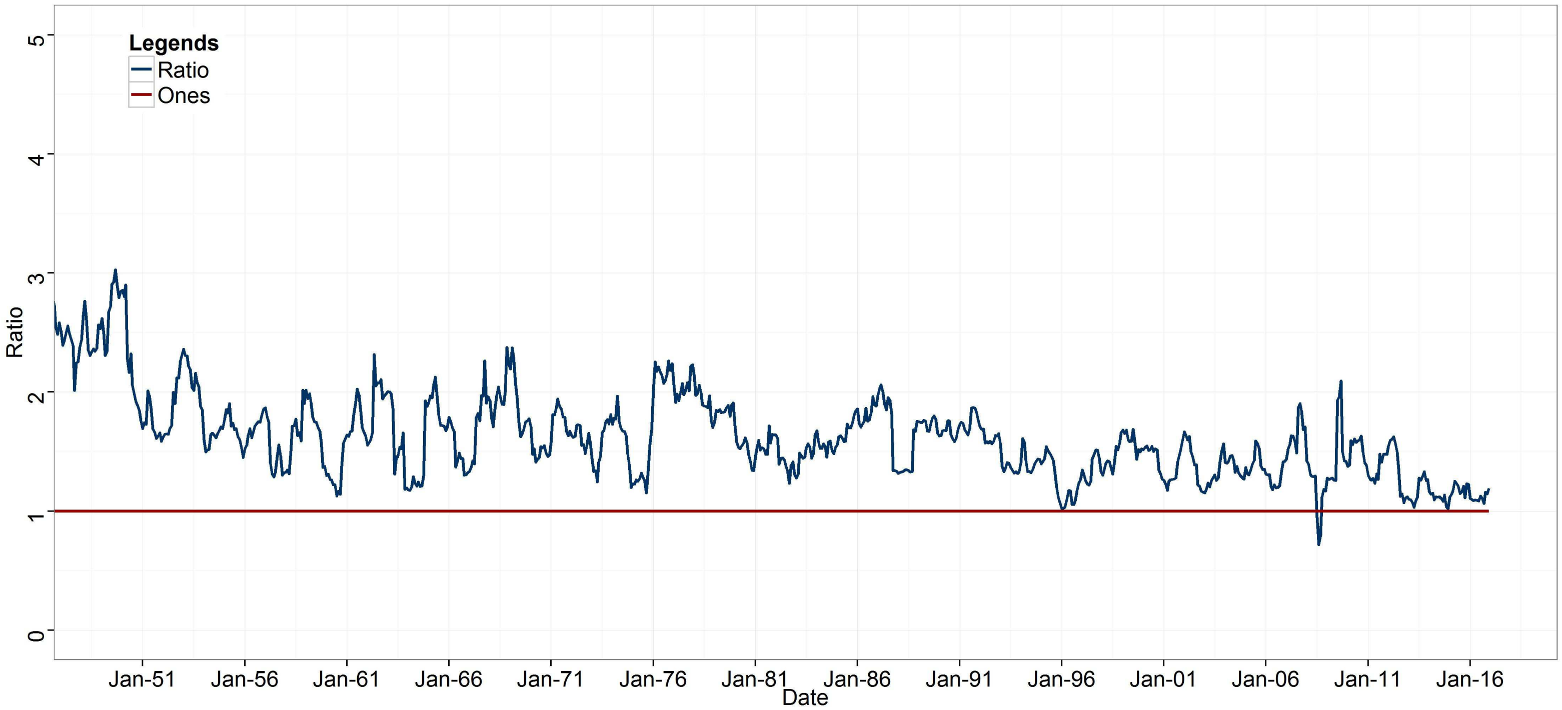

Figure shows the time series of ratio of median sample kurtosis of idiosyncratic daily stock returns for stocks in Group 1 (least actively traded stocks) to the one of Group 10 (most actively traded stocks). The horizontal line marks the value of one. The sample ranges from January 1950 to December 2016.

Figure 3: Kurtosis of Daily Stock Returns for 1950 through 2016.

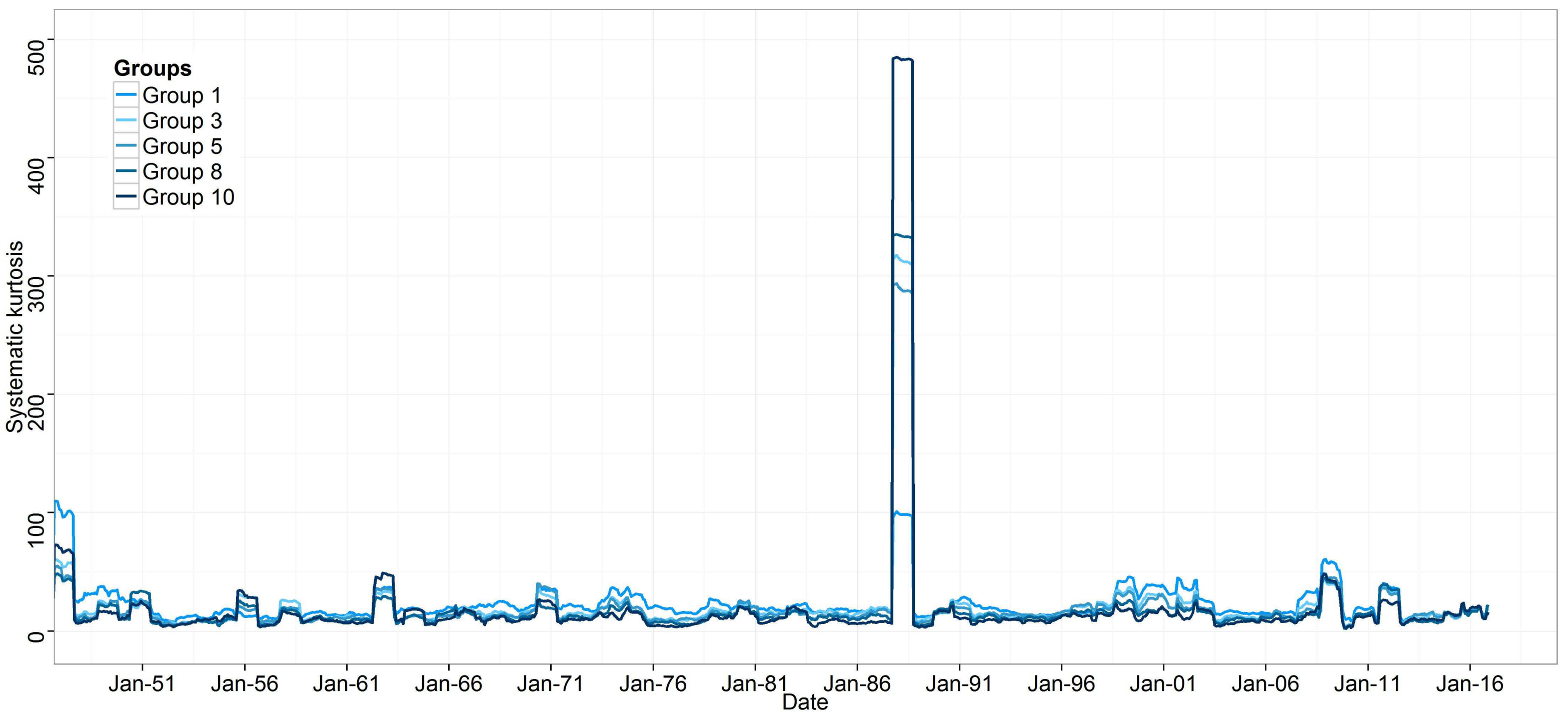

Figure shows five monthly time series of 12-month moving average of median sample kurtosis of daily stock returns for each of the five trading activity groups (groups 1, 3, 5, 8, and 10 out of ten groups). Group 1 (10) contains the least (most) actively traded stocks. The period ranges from January 1950 to December 2016.

Figure 4: Ratio of Kurtosis (Group 1 to Group 10) for 1950 through 2016.

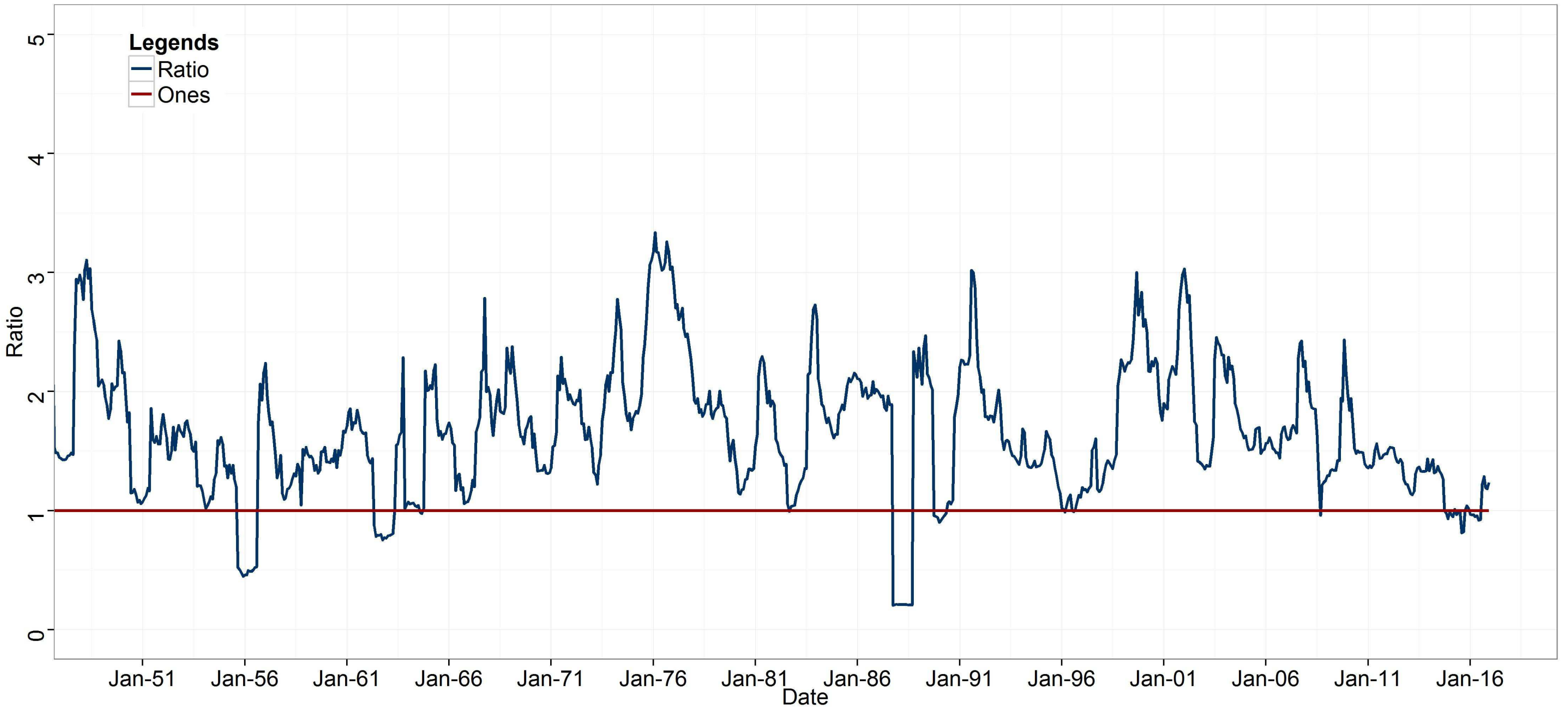

Figure shows the time series of ratio of median sample kurtosis of daily stock returns for stocks in Group 1 (least actively traded stocks) to the one of Group 10 (most actively traded stocks). The horizontal line marks the value of one. The sample ranges from January 1950 to December 2016.

Figure 5: Idiosyncratic Skewness of Daily Stock Returns for 1950 through 2016.

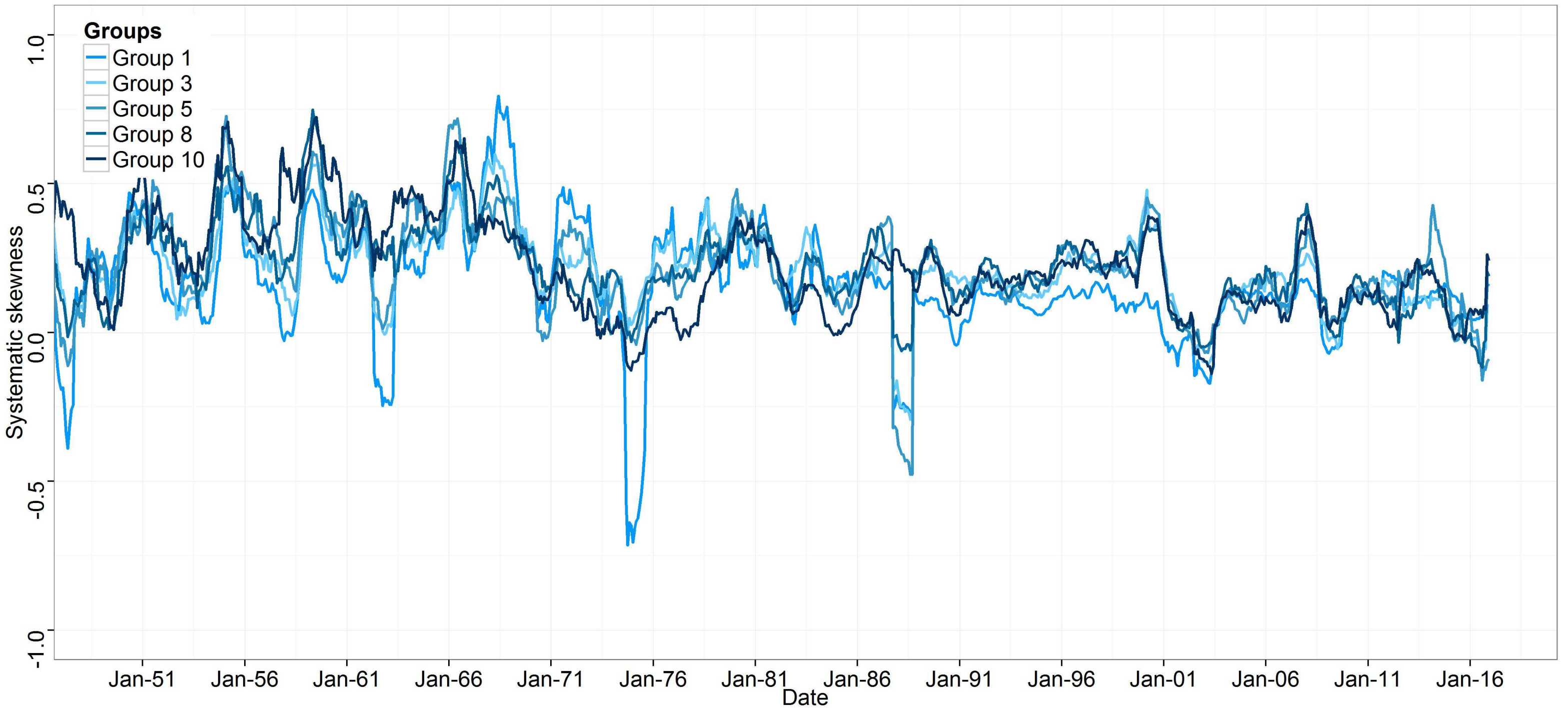

Figure shows five monthly time series of 12-month moving average of median sample skewness of idiosyncratic daily stock returns for each of the five trading activity groups (groups 1, 3, 5, 8, and 10 out of ten groups). Group 1 (10) contains the least (most) actively traded stocks. The period ranges from January 1950 to December 2016.

Figure 6: Idiosyncratic Volatility of Daily Stock Returns for 1950 through 2016.

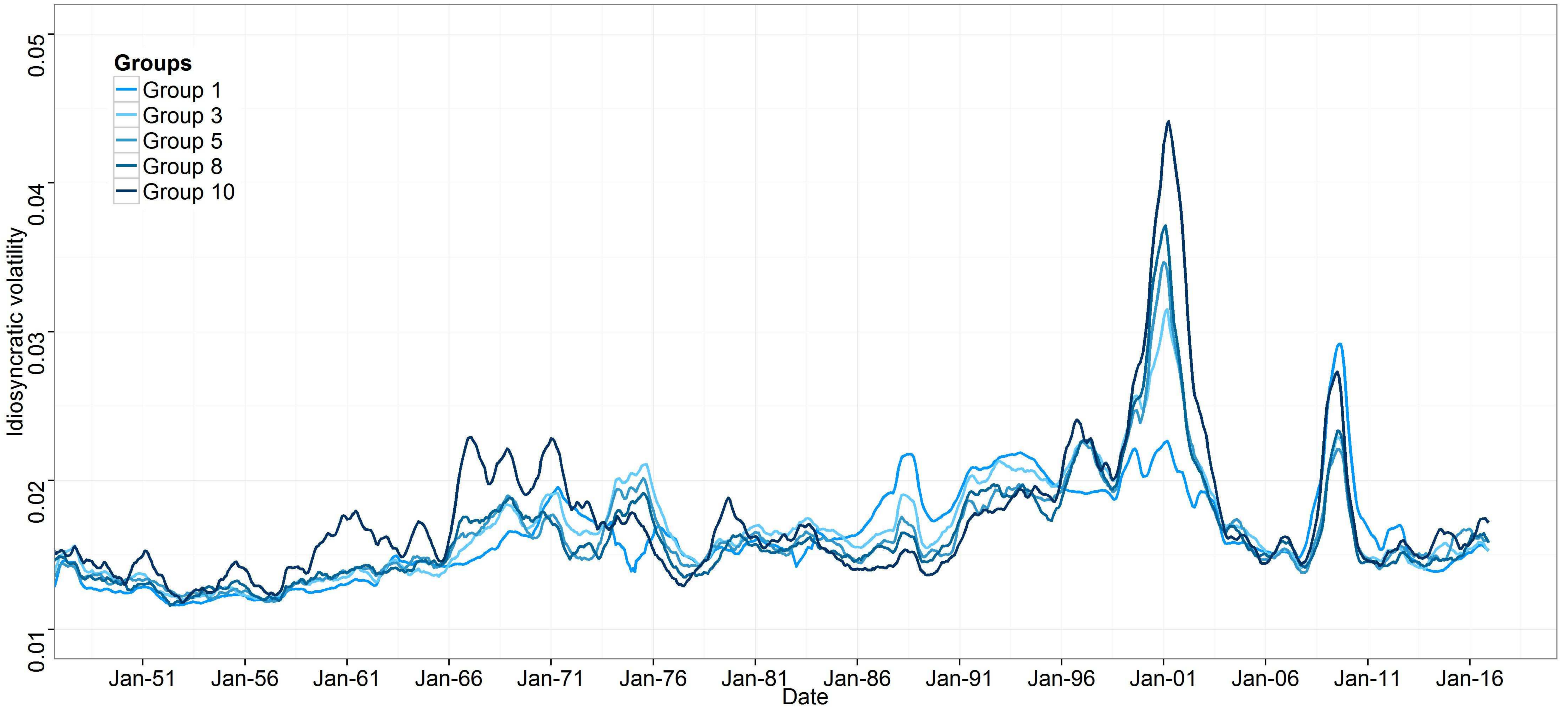

Figure shows five monthly time series of 12-month moving average of median sample volatility of idiosyncratic daily stock returns for each of the five trading activity groups (groups 1, 3, 5, 8, and 10 out of ten groups). Group 1 (10) contains the least (most) actively traded stocks. The period ranges from January 1950 to December 2016.

Table 1: Kurtosis and Skewness for Trading activity Groups across Decades

	Group 1	Group 3	Group 5	Group 8	Group 10	Total
			Decade 1950-1960			
Activity	101	517	967	1,926	5,925	440
Kurtosis	7.174	4.253	4.387	3.205	2.656	4.825
Skewness	0.199	0.243	0.305	0.317	0.357	0.262
# Stocks	676	766	623	511	328	
			Decade 1960-1970			
Activity	277	1,650	3,375	6,999	26,727	1,197
Kurtosis	8.472	5.419	5.188	4.141	2.972	6.002
Skewness	0.303	0.306	0.303	0.302	0.338	0.310
# Stocks	2,126	1,839	1,378	1,056	557	
			Decade 1970-1980			
Activity	210	2,249	5,103	11,519	39,913	1,284
Kurtosis	6.605	5.511	4.928	4.327	3.289	5.627
Skewness	0.182	0.210	0.178	0.159	0.076	0.179
# Stocks	3,697	1,894	1,435	1,028	583	
			Decade 1980-1990			
Activity	1,176	12,893	29,588	70,470	271,798	4,935
Kurtosis	7.133	5.734	4.939	4.321	2.851	5.938
Skewness	0.144	0.201	0.192	0.205	0.157	0.173
# Stocks	5,642	3,084	1,852	1,236	677	
			Decade 1990-2000			
Activity	2,895	38,104	88,975	245,021	1,232,159	14,012
Kurtosis	8.102	6.450	6.174	5.226	4.220	6.884
Skewness	0.098	0.184	0.202	0.197	0.192	0.143
# Stocks	7,721	5,082	3,265	2,237	1,089	
			Decade 2000-2010			
Activity	8,640	190,720	468,578	1,364,880	6,704,911	62,135
Kurtosis	7.583	5.963	5.491	4.699	4.002	6.542
Skewness	0.059	0.098	0.087	0.106	0.096	0.076
# Stocks	5,894	3,545	2,398	1,736	820	
			Decade 2010-2017			
Activity	18,363	411,722	940,703	2,216,188	7,141,896	83,929
Kurtosis	6.808	6.315	6.085	4.918	4.232	6.396
Skewness	0.107	0.074	0.075	0.077	0.098	0.096
# Stocks	3,364	1,659	1,018	766	441	

Table presents the sample medians of trading activity, idiosyncratic skewness, and idiosyncratic kurtosis as well as the number of stocks for the ten groups of U.S. stocks, based on their trading activity. The sample ranges from January 1950 to December 2016 and split into decades. Group 1 (10) consists of stocks with lowest (highest) trading activity in the previous three months.

Table 2: Simulated Theoretical Kurtosis and Low Bounds.

	Group 1	Group 3	Group 5	Group 8	Group 10
Trading Activity	8,000	210,000	460,000	1,000,000	3,600,000
Number of Bets	4	35	59	99	232
Avg Kurtosis	7,214	651	381	225	95
Stand. Error	(4.81)	(0.12)	(0.06)	(0.02)	(0.01)
Low Bound	5,576	631	374	259	95
%Δ	29%	3%	2%	1%	0%

Table reports trading activity $\sigma \cdot V \cdot P$, bet arrival rate per day γ, the average daily returns kurtosis and its standard errors of the means from Monte-carlo simulations, low bound for kurtosis, and percentage difference between the average kurtosis and the low bound for the median stock in each of the five trading activity groups (groups 1, 3, 5, 8, and 10 out of ten groups). Group 1 (10) contains the least (most) actively traded stocks.

Table 3: Imbalance Forecasting Model of Arbitrageurs.

Shredding	const	$S_{i,t-1}$	$S^2_{i,t-1}$	$S_{i,t-2}$	$S^2_{i,t-2}$	$S_{i,t-3}$	$S^2_{i,t-3}$	$S_{i,t-4}$	$S^2_{i,t-4}$	$S_{i,t-5}$	$S^2_{i,t-5}$	R^2
$\eta = 1\%$	39,719	1.98	0.00	0.97	0.00	1.02	0.00	1.25	0.00	5.32	0.00	12%
$\eta = 5\%$	32,925	0.37	0.00	0.29	0.00	0.34	0.00	0.48	0.00	1.49	0.00	13%
$\eta = 10\%$	12,190	0.17	-0.00	0.16	-0.00	0.19	-0.00	0.30	0.00	0.89	-0.00	13%

Table reports estimates $\hat{\beta}_{1j}$ and $\hat{\beta}_{2j},\ j = 1,.5$ of arbitrageurs' model for forecasting unexecuted imbalances

$$S^*_{it,n} - S_{it,n} = \alpha + \sum_{j=1}^{5} \beta_{1j} \cdot S_{i,t-j,n} + \sum_{j=1}^{5} \beta_{2j} \cdot S^2_{i,t-j,n} + \epsilon_{tn}, \qquad t = 1,..T, n = 1,..N,$$

estimated based on the simulated sample for a benchmark stock with daily volatility 2 percent, price $40, and daily volume 1 million shares. The simulated sample consist of 90-day paths. The three bet-shredding algorithms are used: "Method-$V(1\%)$", "Method-$V(5\%)$", and "Method-$V(10\%)$".

Figure 7: Returns Autocorrelations without arbitrager.

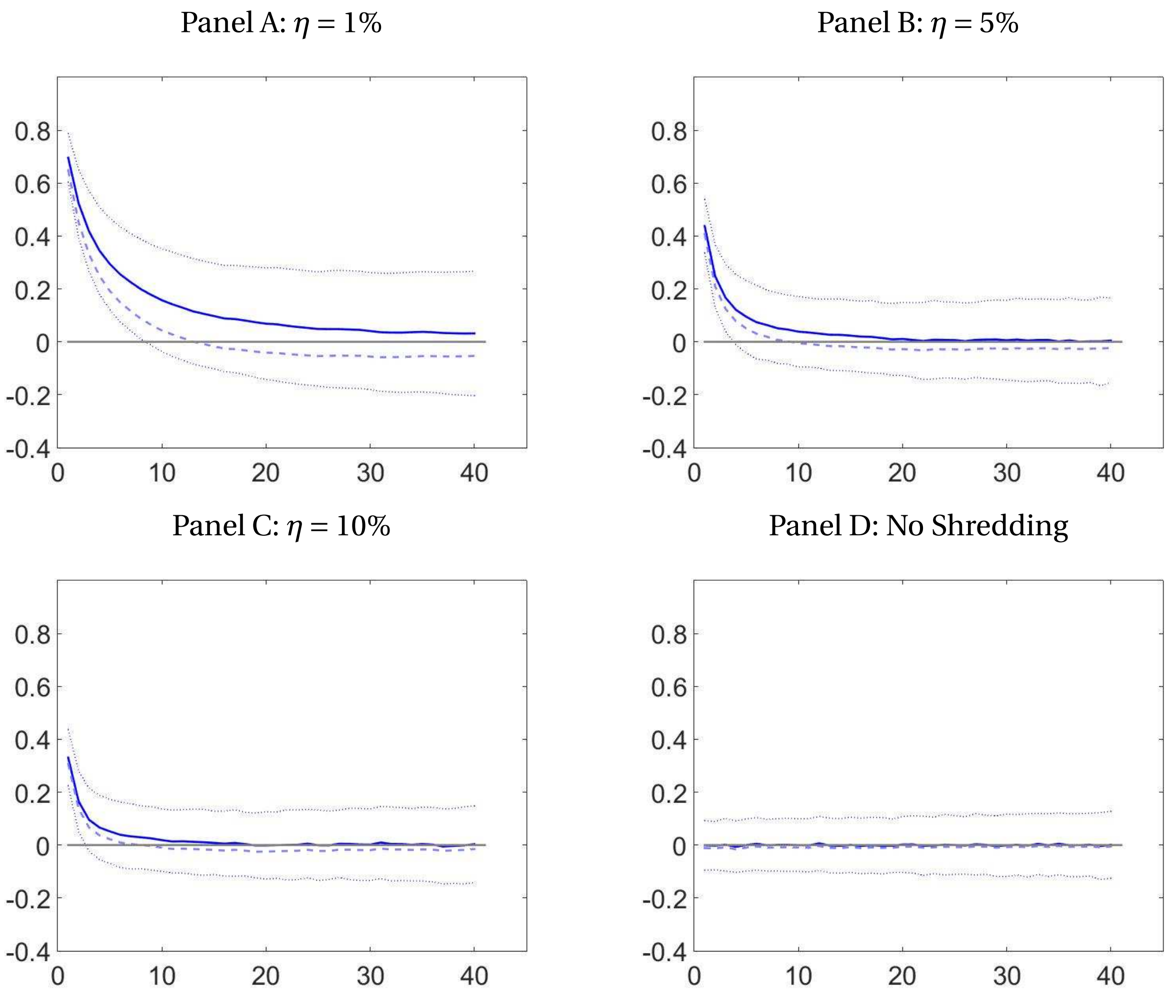

Figure shows autocorrelation coefficients of daily returns at different lags for different models of bet shredding without arbitragers: "Method-$V(1\%)$", "Method-$V(5\%)$", "Method-$V(10\%)$", and no bet shredding. The simulation consists of 90-day paths. There are averages, medians, and standard errors of autocorrelation coefficients in dark solid, dashed, and light solid lines, respectively.

Figure 8: Returns Autocorrelations with arbitrager.

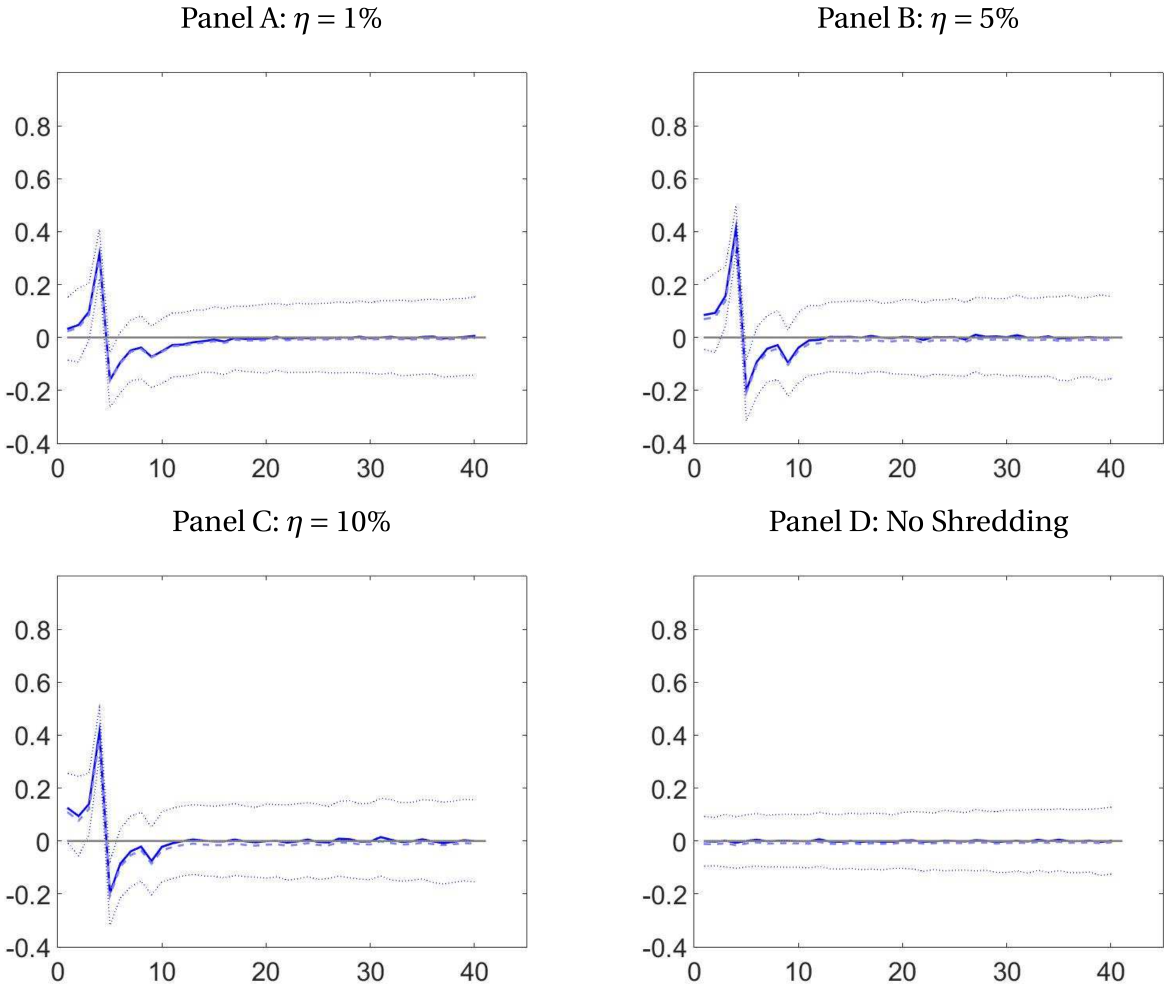

Figure shows average autocorrelation coefficients of daily returns at different lags for different models of bet shredding with arbitragers: “Method-V(1%)”, “Method-V(5%)”, “Method-V(10%)”, and no bet shredding. The simulation consists of 90-day paths. There are averages, medians, and standard errors of autocorrelation coefficients in dark solid, dashed, and light solid lines, respectively.

Table 4: Returns Autocorrelations.

	Order of autocorrelation					
	lag 1	lag 2	lag 3	lag 5	lag 10	lag 20
Panel A: Model without arbitragers						
$\eta = 1\%$	0.696	0.523	0.417	0.294	0.157	0.068
	(0.092)	(0.128)	(0.152)	(0.172)	(0.194)	(0.211)
$\eta = 5\%$	0.437	0.249	0.167	0.094	0.038	0.009
	(0.101)	(0.119)	(0.127)	(0.137)	(0.133)	(0.138)
$\eta = 10\%$	0.331	0.125	0.096	0.051	0.025	-0.001
	(0.106)	(0.114)	(0.119)	(0.121)	(0.118)	(0.127)
No Shredding	-0.001	-0.002	0.001	0.000	0.000	0.000
	(0.093)	(0.091)	(0.099)	(0.099)	(0.099)	(0.10)
Panel B: Model with arbitragers						
$\eta = 1\%$	0.033	0.047	0.098	-0.16	-0.052	-0.005
	(0.119)	(0.140)	(0.106)	(0.103)	(0.122)	(0.130)
$\eta = 5\%$	0.085	0.093	0.157	0.41	-0.04	0.002
	(0.131)	(0.150)	(0.111)	(0.114)	(0.133)	(0.141)
$\eta = 10\%$	0.123	0.094	0.14	0.413	-0.022	0.000
	(0.132)	(0.150)	(0.115)	(0.120)	(0.132)	(0.140)
No Shredding	-0.001	-0.002	0.001	0.000	0.000	0.000
	(0.093)	(0.091)	(0.099)	(0.099)	(0.099)	(0.10)

Table reports average autocorrelation coefficients of daily returns at different lags for different models of bet shredding: “Method-$V(1\%)$”, “Method-$V(5\%)$”, “Method-$V(10\%)$”, and no bet shredding. Panel A presents results for the model without arbitragers. Panel B presents results for the model with arbitragers. The simulation consists of 90-day paths.

Figure 9: Distributions of simulated moments without arbitrageurs

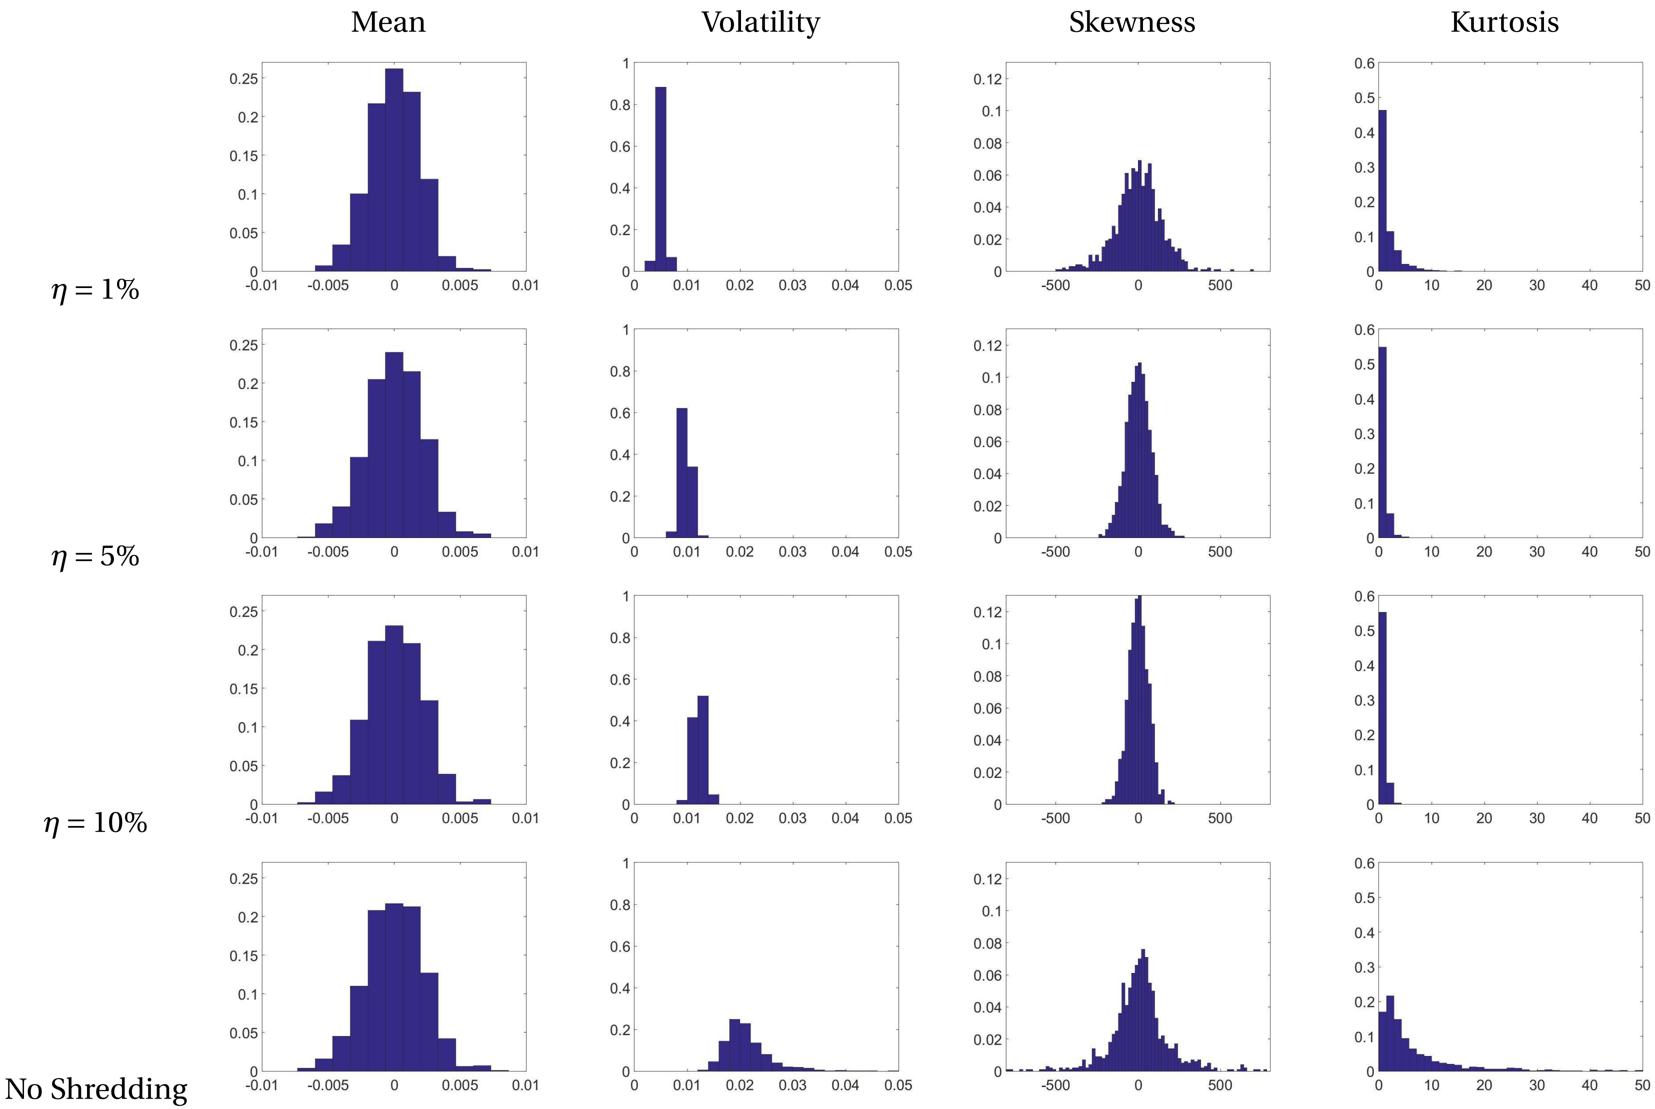

Figure shows distributions of simulated moments of order flow for a benchmark stock. There are 1000 simulations of 90-day paths of returns. The case with no bet shredding and the three bet-shredding algorithms are used: “Method-V(1%)”, “Method-V(5%)”, and “Method-V(10%)”.

Figure 10: Distributions of simulated moments with arbitrageurs

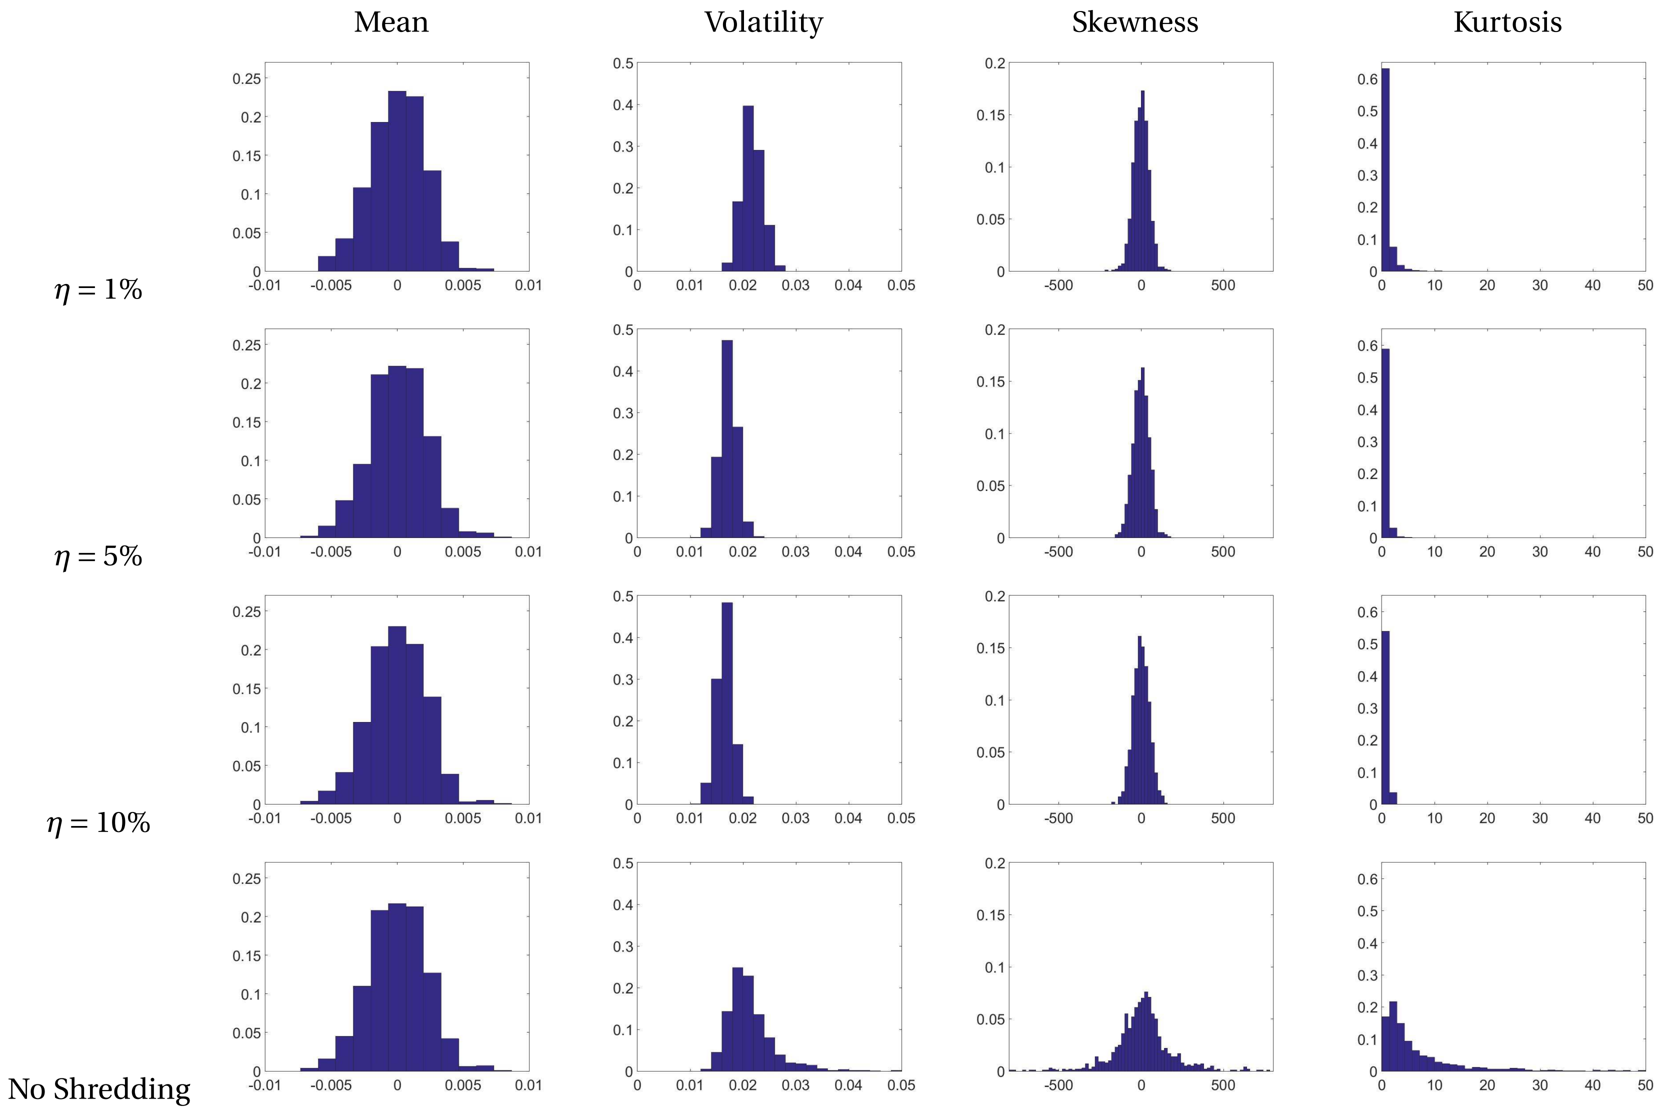

Figure shows distributions of simulated moments of order flow for a benchmark stock with arbitrageurs. There are simulations of 90-day paths of returns. The case with no bet shredding and the three bet-shredding algorithms are used: "Method-V(1%)", "Method-V(5%)", and "Method-V(10%)".

Table 5: Summary Statistics for Daily Returns.

	$\eta = 1\%$	$\eta = 5\%$	$\eta = 10\%$	No Shredding
PANEL A: MODEL WITHOUT ARBITRAGER				
Mean	0.000	0.000	0.000	0.000
	(0.002)	(0.002)	(0.002)	(0.002)
St.dev	0.005	0.010	0.012	0.021
	(0.001)	(0.001)	(0.001)	(0.005)
Skewness	0.252	0.722	0.451	8.564
	(140)	(74)	(61)	(178)
Kurtosis	0.927	0.338	0.281	7.331
	(1.745)	(0.632)	(0.556)	(10.421)
PANEL B: MODEL WITH ARBITRAGER				
Mean	0.000	0.000	0.000	0.000
	(0.002)	(0.002)	(0.002)	(0.002)
St.dev	0.022	0.017	0.017	0.021
	(0.002)	(0.002)	(0.002)	(0.005)
Skewness	-1.152	-0.497	-0.457	8.564
	(46)	(48)	(49)	(178)
Kurtosis	0.517	0.229	0.171	7.331
	(0.943)	(0.553)	(0.556)	(10.421)

Table reports statistics for simulated daily returns such as mean, standard deviation, skewness, and kurtosis for different models of bet shredding: "Method-$V(1\%)$", "Method-$V(5\%)$", "Method-$V(10\%)$", and no bet shredding. Panel A presents results for the model without arbitragers. Panel B presents results for the model with arbitragers. The simulation consists of 90-day paths.

Table 6: Calibrated Bet-Shredding Parameters.

	Group 1	Group 3	Group 5	Group 8	Group 10
			Decade 1950-1960		
W	101	517	967	1,926	5,925
$\hat{\eta}$	8.875	6.288	6.500	3.675	3.225
			Decade 1960-1970		
W	277	1,650	3,375	6,999	26,727
$\hat{\eta}$	9.400	6.888	5.200	4.013	2.588
			Decade 1970-1980		
W	210	2,249	5,103	11,519	39,913
$\hat{\eta}$	9.40	6.89	5.20	4.01	2.59
			Decade 1980-1990		
W	1,176	12,893	29,588	70,470	271,798
$\hat{\eta}$	9.17	4.84	3.71	2.79	1.59
			Decade 1990-2000		
W	2,895	38,104	88,975	245,021	1,232,159
$\hat{\eta}$	8.36	4.22	3.46	2.52	1.66
			Decade 2000-2010		
W	8,640	190,720	468,578	1,364,880	6,704,911
$\hat{\eta}$	6.34	2.94	2.41	1.79	1.37
			Decade 2010-2017		
W	18,363	411,722	940,703	2,216,188	7,141,896
$\hat{\eta}$	5.04	2.69	2.32	1.80	1.39

Table presents calibrated parameter η and trading activity W for the median stocks in the five trading activity groups and for each decade for the period 1950 through 2017.

Table 7: Implied Execution Horizons.

	Group 1	Group 3	Group 5	Group 8	Group 10
			Panel A: Decade 1990-2000		
W	2,895	38,104	88,975	245,021	1,232,159
$\hat{\eta}$	8.36	4.22	3.46	2.52	1.66
std-1	0.02	0.01	0.01	0.00	0.00
std-2	0.11	0.04	0.02	0.02	0.01
std-3	0.55	0.18	0.12	0.08	0.04
std-4	2.69	0.91	0.60	0.39	0.17
std-5	13.22	4.44	2.93	1.92	0.85
std-6	64.86	21.80	14.36	9.42	4.16
std-7	318.23	106.96	70.46	46.22	20.39
			Panel B: Decade 2010-2017		
W	18,363	411,722	940,703	2,216,188	7,141,896
$\hat{\eta}$	5.04	2.69	2.32	1.80	1.39
std-1	0.01	0.00	0.00	0.00	0.00
std-2	0.03	0.01	0.01	0.00	0.00
std-3	0.16	0.04	0.03	0.02	0.01
std-4	0.79	0.19	0.12	0.09	0.05
std-5	3.86	0.91	0.61	0.44	0.26
std-6	18.93	4.46	2.98	2.17	1.29
std-7	92.87	21.88	14.63	10.65	6.32

Table presents implied execution horizons for bets of different sizes for different trading activity groups and time periods. There are calibrated parameter η, trading activity W, and execution horizons (in minutes) for 1 through 7 standard deviation bets.