

Supplementary Material for ‘The Predictive Power of Zero Intelligence in Financial Markets’

J. Doyne Farmer,^{1,*} Paolo Patelli,^{1,2} and Ilija I. Zovko^{1,3}

¹*Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA*

²*Sant’Anna School of Advanced Studies, Piazza Martiri della Libertà 33, Pisa, Italy*

³*CENDEF, University of Amsterdam, Roetersstraat 11, Amsterdam, The Netherlands*

Contents

I. Literature review	1
A. Standard literature	1
B. Random process models of continuous double auction	2
II. Dimensional analysis	3
III. The London Stock Exchange (LSE) data set	3
A. Opening auction, real order types, time	4
IV. Measurement of model parameters	4
A. Measuring μ and σ	5
B. Measuring α and δ	5
C. Measuring the price diffusion rate	6
V. Estimating the errors for the regressions	6
VI. Market impact	7
A. Relation of market impact to supply and demand schedules	7
B. Alternative market impact collapse plots	7
C. Error analysis for market impact	7
D. Market impact in log-log coordinates	10
VII. Extending the model	11
References and Notes	12

I. LITERATURE REVIEW

We will begin by reviewing what we will call the “standard literature”, which includes both purely empirical studies, as well as theoretical studies predicated on rationality models. We will then review the literature on random process models of the continuous double auction, which is more closely related to the model we test here.

A. Standard literature

The market microstructure literature focusing on the understanding of spread, volatility and market impact in financial markets is theoretically and empirically extensive. The theoretical analyses traditionally use the underlying paradigm of rational agents. Models of spread, starting with Demsetz (1968), Tinic (1972), Stoll (1978), Amihud and Mendelson (1980), Ho and Stoll (1981) [1, 22, 43, 61, 64], have examined the possible determinants of spreads as a result of rational, utility-maximizing problem faced by the market makers. Models providing insight into the utility-maximizing response of agents to other various measures of market conditions such as volatility are for example Lo, MacKinlay, and Zhang (2002) [48] who investigate a simple model in which the log stock price is modeled as a Brownian motion diffusion process. Provided agents prefer a lower expected execution time, their model predicts a positive relationship between volatility and limit order placement. Copeland and Galai (1983), Glosten and Milgrom (1985), Easley and OHara (1987), Glosten (1995), Foucault (1999), and Easley, OHara, and Saar (2001) [20, 24, 25, 31, 38, 39] examine asymmetric information effects on order placement. Andersen (1996) [2] modifies the Glosten and Milgrom (1985) [39] model with the stochastic volatility and information flow perspective. Other models of trading in limit order markets include Cohen et al. (1981), Angel (1994), Harris (1998), Chakravarty and Holden (1995), Seppi (1997) and Rock (1990), Parlour and Seppi (2001), Parlour (1998), Foucault, Kadan, and Kandel (2001), Domowitz and Wang (1994) [3, 13, 19, 23, 32, 40, 53, 54, 56, 58].

Empirical research is equally rich. Roll (1984) and Choi et al (1988) [16, 57] estimate spreads from transaction prices. Glosten and Harris (1988), Hasbrouck (1988, 1991), Madhavan, Richardson and Roomans (1996) [37, 41, 50] model spreads as VAR using various trade variables as independent variables. Roll (1984), Choi, Salandro, Shastri (1988), George, Kaul and Nimalendran (1991), Huang and Stoll (1997), Jang and Venkatesh (1991) [16, 36, 44, 45, 57] focus on estimating components of the spread.

Empirical research in volatility was initiated by statistical descriptions of the volatility process (Engle (1982), Bollerslev (1986) [8, 27], Bollerslev, Chou, and Kroner (1992) [9] for a survey), but has grown increasingly ambitious with multivariate structural models of the inter-

*Electronic address: jdf@santafe.edu

action between volatility and other economic variables. Positive correlation of daily volume and volatility has been documented in Clark (1973), Epps and Epps (1976), Tauchen and Pitts (1983) [17, 28, 63]. Volume has been entered into ARCH specifications by Lamoureux and Lastrapes (1994) [46]. Other empirical investigations of volatility determinants and consequences include Gallant et al (1992), Anderson (1996), Blume et al (1991), Reiss and Werner (1994), Fleming et al (2001), Hasbrouck and Saar (2003) [2, 7, 30, 35, 42, 55]. A comprehensive study of the joint distribution of returns and volume is done by Gallant, Rossi, and Tauchen (1992, 1993) [34, 35].

B. Random process models of continuous double auction

There are two independent lines of prior work, one in the financial economics literature, and the other in the physics literature. The models in the economics literature are directed toward econometrics and treat the order process as static. In contrast, the models in the physics literature are mostly conceptual toy models, but they allow the order process to react to changes in prices, and are thus fully dynamic. Our model bridges this gap. This is explained in more detail below.

The first model of this type that we are aware of in the economics literature was due to Mendelson [52], who modeled random order placement with periodic clearing. Cohen *et al.* [18] developed a model of a continuous auction, modeling limit orders, market orders, and order cancellation as Poisson processes. However, they only allowed limit orders at two fixed prices, buy orders at the best bid, and sell orders at the best ask. This assumption allowed them to use standard results from queuing theory to compute properties such as the expected number of stored limit orders, the expected time to execution, and the relative probability of execution vs. cancellation. Domowitz and Wang [23] extended this to multiple price levels by assuming arbitrary order placement and cancellation processes (which can take on any value at each price level). They assume that these processes are fixed in time, and do not respond to changes in the best bid or ask. This allows them to derive the distribution of the spread, transaction prices, and waiting times for execution. This model was tested by Bollerslev *et al.* [10] on three weeks of data for the Deutschemark/U.S. Dollar exchange rate. They showed that it does a good job of predicting the distribution of the spread. However, since the prices are pinned, the model does not make a prediction about price diffusion, and this also creates errors in the predictions of the spread and stored supply and demand.

The models in the physics literature, which appear to have been developed independently, differ in that they address price dynamics. That is, they incorporate the feedback between order placement and price formation, allowing the order placement process to change in re-

sponse to changes in prices. These models have mainly been conceptual toy models designed to understand the anomalous diffusion properties of prices (a property that all of these models fail to reproduce, as explained later). This line of work begins with a paper by Bak *et al.* [4] which was developed by Eliezer and Kogan [26] and by Tang [62]. They assume that limit orders are placed at a fixed distance from the midpoint, and that the limit prices of these orders are then randomly shuffled until they result in transactions. It is the random shuffling that causes price diffusion. This assumption, which we feel is unrealistic, was made to take advantage of the analogy to a standard reaction-diffusion model in the physics literature. Maslov [51] introduced an alternative model that was solved analytically in the mean-field limit by Slanina [59]. Each order is randomly chosen to be either a buy or a sell with equal probability, and either a limit order or a market order with equal probability. If a limit order, it is randomly placed within a fixed distance of the current price. Both the Bak et al. model and that of Maslov result in anomalous price diffusion, in the sense that the Hurst exponent $H = 1/4$ (in contrast to standard diffusion, which has $H = 1/2$, or real prices which tend to have $H > 1/2$). In addition, the Maslov model unrealistically requires equal probabilities for limit and market order placement, otherwise the inventory of stored limit orders either goes to zero or grows without bound. A model adding a Poisson order cancellation process was proposed by Challet and Stinchcombe [14], and independently by Daniels et al. [21]. Challet and Stinchcombe showed that this results in $H = 1/4$ for short times, but asymptotically gives $H = 1/2$. The Challet and Stinchcombe model, which posits an arbitrary, unspecified function for the relative position of limit order placement, is quite similar to that of Domowitz and Wang [23], but allows for the possibility of order placement responding to price movement.

The model we test here was introduced by Daniels *et al.* [21]. Like other physics models, it treats the feedback between order placement and price movement. It has the advantage that it is defined in terms of five scalar parameters, and so is parsimonious and can easily be tested against real data. Its simplicity enables a dimensional analysis, which gives approximate predictions about many of the properties of the model. Perhaps most important is the use to which the model is put: With the exception of reference [26], work in the physics literature has focused almost entirely on the anomalous diffusion of prices. While interesting and important for refining risk calculations, from a practical point of view this is a second-order effect. In contrast, the model studied here focuses on first order effects of primary interest to market participants, such as the bid-ask spread, volatility, depth profile, price impact, and the probability and time to fill an order. It demonstrates how dimensional analysis becomes a useful tool in an economic setting, and the analysis done in Daniels *et al.* and Smith *et al.* develops mean field theories to understand many relevant mar-

ket properties. Many of the important properties of the model can be stated in terms of simple scaling relations in terms of the five parameters.

Subsequent to reference [21], Bouchaud et al. [12] demonstrated that they can derive a simple equation for the depth profile, by making the assumption that prices execute a random walk and introducing an additional free parameter. In this paper we show how to do this from first principles without introducing a free parameter. Iori and Chiarella [15] have numerically studied fundamentalists and technical traders placing limit orders; a talk on this work by Giulia Iori in part inspired this model.

II. DIMENSIONAL ANALYSIS

Dimensional analysis can be used to simplify the study of this model and to make some approximate predictions about several of its properties. For a good reference on dimensional analysis see Barenblatt [5].

There are three fundamental dimensional quantities in this model: *shares*, *price*, and *time*. There are five parameters. Because they have independent dimensions, when the dimensional constraints between the parameters are taken into account, this leaves only two independent degrees of freedom. It turns out that the order flow rates μ , α , and δ are more important than the discreteness parameters σ and dp , in the sense that the properties of the model are much more sensitive to variations in the order flow rates than they are to variations in σ or dp . It therefore natural to construct non-dimensional units based on the order flow parameters alone. There are unique combinations of the three order flow rates with units of *shares*, *price*, and *time*. This gives characteristic scales for price, shares, and time, that are unique up to a constant: the characteristic number of shares $N_c = \mu/\delta$, the characteristic price interval $p_c = \mu/\alpha$, and the characteristic timescale $t_c = 1/\delta$. These are the unique combinations of these three parameters with the correct dimensions.

These characteristic scales can be used to define non-dimensional coordinates based on the order flow rates. These are $\hat{p} = p/p_c$ for price, $\hat{N} = N/N_c$ for shares, and $\hat{t} = t/t_c$ for time. The use of non-dimensional coordinates has the great advantage that it reduces the number of degrees of freedom from five to two. That is, instead of five independent parameters, we only have two independent parameters, which makes the model easier to understand.

The two irreducible degrees of freedom are naturally discussed in terms of non-dimensional versions of the discreteness parameters. A non-dimensional scale parameter based on order size is constructed by dividing the typical order size σ (with dimensions of *shares*) by the characteristic number of shares N_c . This gives the non-dimensional parameter $\epsilon \equiv \sigma/N_c = \delta\sigma/\mu$, which characterizes the granularity of the order flow. A non-dimensional scale parameter based on tick size is constructed by dividing the tick size dp by the characteristic price, i.e. $dp/p_c = \alpha dp/\mu$. The usefulness of this is that

the properties of the model only depend on the two non-dimensional parameters, ϵ and dp/p_c : Any variations of the parameters μ , α , and δ that keep these two non-dimensional parameters constant gives exactly the same market properties. One of the interesting results that emerges from analysis of the model is that the effect of the granularity parameter ϵ is generally much more important than the tick size dp/p_c . For a more detailed discussion, see reference [60].

While we have investigated numerically the effect of varying the tick size in ref. [60], for the purposes of comparing to data, here we simply take the limit $dp \rightarrow 0$, which provides a reasonable approximation.

III. THE LONDON STOCK EXCHANGE (LSE) DATA SET

The London Stock Exchange is composed of two parts, the electronic open limit order book, and the upstairs quotation market, which is used to facilitate large block trades. During the time period of our dataset 40% to 50% of total volume was routed through the electronic order book and the rest through the upstairs market. It is believed that the limit order book is the dominant price formation mechanism of the London Stock Exchange: about 75% of upstairs trades happen between the current best prices in the order book [49]. Our analysis involves only the data from the electronic order book. We chose to study this data set because we have a complete record of every action taken by every participating institution, allowing us to measure the order flows and cancellations and estimate all of the necessary parameters of our model.

We used data from the time period August 1st 1998 - April 30th 2000, which includes a total of 434 trading days and roughly six million events. We chose 11 stocks each having the property that the number of total number of events exceeds 300,000 and was never less than 80 on any given day. Some statistics about the order flow for each stock are given in table I.

The trading day of the LSE starts at 7:50 with a roughly 10 minute long opening auction period (during the later part of the dataset the auction end time varies randomly by 30 seconds). During this time orders accumulate without transactions; then a clearing price for the opening auction is calculated, and all opening transactions take place at this price. Following the opening at 8:00 the market runs continuously, with orders matched according to price and time priority, until the market closes at 16:30. In the earlier part of the dataset, until September 22nd 1999, the market opening hour was 9:00. During the period we study there have been some minor modifications of the opening auction mechanism, but since we discard the opening auction data anyway this is not relevant.

Some stocks in our sample (VOD for example) have had stock price splits and tick price changes during the

stock ticker	num. events (1000s)	average (per day)	limit (1000s)	market (1000s)	deletions (1000s)	eff. limit (shares)	eff. market (shares)	# days
AZN	608	1405	292	128	188	4,967	4,921	429
BARC	571	1318	271	128	172	7,370	6,406	433
CW.	511	1184	244	134	134	12,671	11,151	432
GLXO	814	1885	390	200	225	8,927	6,573	434
LLOY	644	1485	302	184	159	13,846	11,376	434
ORA	314	884	153	57	104	12,097	11,690	432
PRU	422	978	201	94	127	9,502	8,597	354
RTR	408	951	195	100	112	16,433	9,965	431
SB.	665	1526	319	176	170	13,589	12,157	426
SHEL	592	1367	277	159	156	44,165	30,133	429
VOD	940	2161	437	296	207	89,550	71,121	434

TABLE I: Summary statistics for stocks in the dataset. Fields from left to right: stock ticker symbol, total number of events (effective market orders + effective limit orders + order cancellations) in thousands, average number of events in a trading day, number of effective limit orders in thousands, number of effective market orders in thousands, number of order deletions in thousands, average limit order size in shares, average market order size in shares, number of trading days in the sample.

period of our sample. We take splits into account by transforming stock sizes and prices to pre-split values. In any case, since all measured quantities are in logarithmic units, of the form $\log(p_1) - \log(p_2)$, the absolute price scale drops out. Our theory predicts that the tick size should change some of the quantities of interest, such as the bid-ask spread, but the predicted changes are small enough in comparison with the effect of other parameters that we simply ignore them (and base our predictions on the limit where the tick size is zero). Since granularity is much more important than tick size, this seems to be a good approximation.

A. Opening auction, real order types, time

Since the model does not take the opening auction into account, we simply neglect orders leading up to the opening auction, and base all our measurements on the remaining part of the trading day, when the auction is continuous.

In order to treat simply and in a unified manner the diverse types of orders traders can submit in a real market (for example, crossing limit orders, market orders with limiting price, ‘fill-or-kill, execute & eliminate) we use redefinitions based on whether an order results in an immediate transaction, in which case we call it an *effective market order*, or whether it leaves a limit order sitting in the book, in which case we call it an *effective limit order*. Marketable limit orders (also called crossing limit orders) are limit orders that cross the opposing best price, and so result in at least a partial transaction. The portion of the order that results in an immediate transaction is counted as an effective market order, while the non-transacted part (if any) is counted as an effective limit order. Orders that do not result in a transaction and do not leave a limit order in the book, such as for example, failed fill-or-kill orders, are ignored altogether. These have no effect on prices, and in any case, make up only a very small

fraction of the order flow, typically less than 1%. Note that we drop the term “effective”, so that e.g. “market order” means “effective market order”.

A limit order can be removed from the book for many reasons, e.g. because the agent changes her mind, because a time specified when the order was placed has been reached, or because of the institutionally-mandated 30 day limit on order duration. We will lump all of these together, and simply refer to them as “cancellations”.

Our measure of time is based on the number of events, i.e., the time elapsed during a given period is just the total number of events, including effective market order placements, effective limit order placements, and cancellations. We call this *event time*. Price intervals are computed as the difference in the logarithm of prices, which is consistent with the model, in which all prices are assumed to be logarithmic in order to assure their positiveness.

IV. MEASUREMENT OF MODEL PARAMETERS

We test the predictions of the model against real data cross-sectionally on eleven stocks. The parameters of the model are stated in terms of order arrival rates, cancellation rate, order size, and our dataset allows us to compute for each stock the average values of these parameters. As we explain here, these average rates are calculated as means of the daily values weighted by the daily number of events. An alternative would have been to calculate the mean values of the parameters over the entire 2 year period for each stock. While this works well for the parameters μ and σ , it does not work as well for α and δ , as explained below.

A. Measuring μ and σ

The parameter μ_t , which characterizes the average market order arrival rate on day t is just the ratio of the number of shares of effective market orders (for both buy and sell orders) to the number of events during the trading day. Thus for μ it makes no difference whether we measure it across the whole period, or take a weighted average of daily values. This is also true for the average order size σ_t . One complication in measuring σ is that the model assumes that the average size for limit orders and market orders is the same, whereas for the real data this is not strictly true. Nonetheless, as seen in Table I, although the limit order size tends to be a bit larger than the market order size, it is still a fairly good approximation to take them to be the same. For the purposes of this analysis we use the limit order size to measure order size, based on theoretical arguments that this is more important than the market order size. In any case, this does not make a significant difference in the results.

B. Measuring α and δ

The measuring of the cancellation rate δ_t and the limit order rate density α_t is more complicated, due to the highly simplified assumptions we made for the model. In contrast to our assumption of a constant density for placement of limit orders across the entire logarithmic price axis, real limit order placement is highly concentrated near the best prices. Roughly 2/3 of all orders are placed either at the best price or inside the spread. Outside the spread the density of limit order placement falls off as a power law as a function of the distance from the best prices [12, 65]. In addition, we have assumed a constant cancellation rate, whereas in reality orders placed near the best prices tend to be cancelled much faster than orders placed far from the best prices. We cope with these problems by introducing an auxiliary assumption. Basically, we assume that order placement is constant inside an interval, and is zero outside that interval. This is described in more detail below.

In order to estimate the limit order rate density for day t , α_t , we make an empirical estimate of the distribution of the relative price for effective limit order placement on each day. For buy orders we define the relative price as $\Delta = m - p$, where p is the logarithm of the limit price and m is the logarithm of the midquote price. Similarly for sell orders, $\Delta = p - m$. We then somewhat arbitrarily choose Q_t^{lower} as the 2 percentile of the density of Δ corresponding to the limit orders arriving on day t , and Q_t^{upper} as the 60 percentile of Δ . Assuming constant density within this range, we calculate α_t as $\alpha_t = L / (Q_t^{\text{upper}} - Q_t^{\text{lower}})$ where L is the total number of shares of effective limit orders within the price interval $(Q_t^{\text{lower}}, Q_t^{\text{upper}})$ on day t . The choice of $(Q_t^{\text{lower}}, Q_t^{\text{upper}})$ is made in a compromise to include as much data as possi-

ble for statistical stability, but not so much as to include orders that are unlikely to ever be executed, and therefore unlikely to have any effect on prices.

Similarly, to cope with the fact that in reality the average cancellation rate δ decreases [12] with the relative price Δ , whereas in the model δ is assumed to be constant, we base our estimate for δ only on canceled limit orders within the range of the same relative price boundaries $(Q_t^{\text{lower}}, Q_t^{\text{upper}})$ defined above. We do this to be consistent in our choice of which orders are assumed to contribute significantly to price formation (orders closer to the best prices contribute more than orders that are further away). We then measure δ_t , the cancellation rate on day t , as the inverse of the average lifetime of a canceled limit order in the above price range. Lifetime is measured in terms of number of events happening between the introduction of the order and its subsequent cancellation.

The parameter Q_t^{upper} is referred to as W in the main text. In other subsequent studies (to be reported elsewhere) we are able to set the parameter $(Q_t^{\text{lower}} = 0$, and to compute Δ relative to the opposite best price rather than the midprice, with negligible differences in the results. The difference is that in the later studies we have a cleaner dataset. In this dataset there are some points that are clearly outliers, and it was convenient to introduce a lower cutoff for outlier removal. Thus, we do not feel that $(Q_t^{\text{lower}}$ is an important parameter for this analysis, and we have not discussed it in the text (where we have limited space).

The use of this procedure dictates that it is better to choose an average of daily parameters rather than computing average parameters based on ratios of values for the whole period. Because the width of the interval over which orders are placed varies significantly in time. Moment-by-moment orderbook reconstruction makes it clear that the properties of the market tend to be relatively stationary during each day, changing more dramatically overnight. The order-flows on different days can be rather dissimilar. This non-stationarity of the order flow means that δ and α parameters calculation would perform poorly if we attempted to use an average price interval over the whole period. This would have the result that on some days we might count only a small fraction of the orderflow, excluding many orders that were important for price formation, while on other days we would include almost all the orders, including many that were not very relevant for price formation. This problem makes it natural to use daily averages of parameters.

This introduces the worry that daily variations in W might be an important predictive variable, above and beyond its effect on changing α (which is consistent with the model). There is a tendency for the value of W on a given day to track the spread, due to regularities in order placement, and therefore to automatically have some correlation with the spread. We have done several studies, which will be reported in a future work, testing the importance of this effect. These show that while daily

variations in W do give additional predictability for the spread, other aspects of the model are substantially responsible for these results.

C. Measuring the price diffusion rate

The measurement of the price diffusion rate requires some discussion. We measure the intraday price diffusion by computing the mid-point price variance $V(\tau) = \text{Var}\{m(t+\tau) - m(t)\}$, for different time scales τ . The averaging over t includes all events that change the mid-point price. The plot of $V(\tau)$ against τ is called a diffusion curve and for an IID random walk is a straight line with slope D , the diffusion coefficient.

In our case, the computation of D is as follows: For each day we compute the diffusion curve. In this way we avoid the overnight price changes which would bias our estimate. To the daily diffusion curve we then fit a straight line $V(\tau) = D_t\tau$ using least squares weighted by the square root of the number of observations for each value of $V(\tau)$. In fitting a straight line we are assuming IID mid-point price movement which is relatively well born out in the data. For an example of see Fig. 1. Averaging the daily diffusion rates we obtain the full sample estimate of the stock diffusion. We weigh the averaging by the number of events in each day. One must bear in mind that the price diffusion rate from day to day has substantial correlations, as illustrated in Fig. 2.




FIG. 1: Illustration of the procedure for measuring the price diffusion rate for Vodafone (VOD) on August 4th, 1998. On the x axis we plot the time τ in units of ticks, and on the y axis the variance of mid-price diffusion $V(\tau)$. According to the hypothesis that mid-price diffusion is an uncorrelated Gaussian random walk, the plot should obey $V(\tau) = D\tau$. To cope with the fact that points with larger values of τ have fewer independent intervals and are less statistically significant, we use a weighted regression to compute the slope D .




FIG. 2: Time series (top) and autocorrelation function (bottom) for daily price diffusion rate D_t for Vodafone. Because of long-memory effects and the short length of the series, the long-lag coefficients are poorly determined; the figure is just to demonstrate that the correlations are quite large.

V. ESTIMATING THE ERRORS FOR THE REGRESSIONS

The error bars presented in the text are based on a bootstrapping method. It may at first seem that the proper method would be to simply use White's heteroskedasticity consistent estimators, however, we are driven to use this method for two reasons.

First, within each stock the daily values of the dependent variables display slowly decaying positive autocorrelation functions. Averaging the daily values to get an estimate of the stock-specific average may seem to remedy the autocorrelation problem. However, the autocorrelation is very persistent, almost to the scale of the length of our dataset, and the variables may indeed have long memory.¹ This makes us suspicious about using standard statistics.

A second reason in using a bootstrapping method for inference is the fact that, in addition to possibly being long memory, the daily values of the variables are cross-correlated across stocks. (A high volatility in one stock on a particular day is likely to be associated with high volatility in other stocks.) These two reasons lead us to believe that using standard or White's estimators would

¹ It has recently been shown that order sign, order volume and liquidity as reflected by volume at the best price, are long-memory processes [11, 47].

underestimate the regression errors.

The method we use is inspired by the variance plot method described in Beran [6], Section 4.4. We divide the sample into blocks, apply the regression to each block, and then study the scaling of the deviation in the results as the blocks are made longer to coincide with the full sample. We divide the N daily data points for each stock into m disjoint blocks, each containing n adjacent days, so that $n \approx N/m$. We use the same partition for each stock, so that corresponding blocks for each stock are contemporaneous. We perform an independent regression on each of the m blocks, and calculate the mean M_m and standard deviation σ_m of the m slope parameters A_i and intercept parameters B_i , $i = 1, \dots, m$. We then vary m and study the scaling as shown in Figs. 3 and 4.

Figs. 3(a) and (b) illustrate this procedure for the spread, and Figs. 4(a) and (b) illustrate this for the price diffusion rate. Similarly, panels (c) and (d) in each figure show the mean and standard deviation for the intercept and slope as a function of the number of bins. As expected, the standard deviations of the estimates decreases as n increases. The logarithm of the standard deviation for the intercept and slope as a function of $\log n$ is shown in panels (e) and (f). For IID normally distributed data we expect a line with slope $\gamma = -1/2$; instead we observe $\gamma > -1/2$. For example for the spread $\gamma \approx -0.19$. $|\gamma| < 1/2$ is an indication that this is a long memory process.

This method can be used to extrapolate the error for $m = 1$, i.e. the full sample. This is illustrated in panels (e) and (f) in each figure. The inaccuracy in these error bars is evident in the unevenness of the scaling. This is particularly true for the price diffusion rate. To get a feeling for the accuracy of the error bars, we estimate the standard deviation for the scaling regression assuming standard error, and repeat the extrapolation for the one standard deviation positive and negative deviations of the regression lines, as shown in panels (e) and (f) of Figs. 3 and 4. The results are summarized in Table II.

One of the effects that is evident in Figs. 3(c-d) and 4(c-d) is that the slope coefficients tend to decrease as m increases. We believe this is due to the autocorrelation bias discussed in Section (V).

VI. MARKET IMPACT

A. Relation of market impact to supply and demand schedules

The market impact function is closely related to the more familiar notions of supply and demand. We have chosen to measure average market impact in this paper rather than average relative supply and demand for reasons of convenience. Measuring the average relative supply and demand requires reconstructing the limit order book at each instant, which is both time consuming and error prone. The average market impact function, in con-

trast, can be measured based on a time series of orders and best bid and ask prices.

At any instant in time the stored queue of sell limit orders reveals the quantity available for sale at each price, thus showing the supply, and the stored buy orders similarly show the revealed demand. The price shift caused by a market order of a given size depends on the stored supply or demand through a moment expansion [60]. Thus, the collapse of the market impact function reflects a corresponding property of supply and demand. Normally one would assume that supply and demand are functions of human production and desire; the results we have presented here suggest that on a short timescale in financial markets their form is dictated by the dynamical interaction of order accumulation, removal by market orders and cancellation, and price diffusion.

B. Alternative market impact collapse plots

We have demonstrated a good collapse of the market impact using nondimensional units. However, in deciding what “good” means, one should compare this to the best alternatives available. We compare to three such alternatives. In figure 5, the top left panel shows the collapse when using non-dimensional units derived from the model (repeated from the main text). The top right panel shows the average market impact when we instead normalise the order size by its sample mean. Order size is measured in units of shares and market impact is in log price difference. The bottom left panel attempts to take into account daily variations of trading volume, normalizing the order size by the average order size for that stock on that day. In the bottom right panel we use trade price to normalize the order sizes which are now in monetary units (British Pounds). We visually see that none of the alternative rescalings comes close to the collapse we obtain when using non-dimensional units; because of the much greater dispersion, the error bars in each case are much larger.

C. Error analysis for market impact

Assigning error bars to the average market impact is difficult because the absolute price changes Δp have a slowly decaying positive autocorrelation function. This may be a long-memory process, although this is not as obvious as it is for other properties of the market, such as the volume and sign of orders [11, 47]. The *signed* price changes Δp have an autocorrelation function that rapidly decays to zero, but to compute market impact we sort the values into bins, and all the values in the bin have the same sign. One might have supposed that because the points entering a given bin are not sequential in time, the correlation would be sufficiently low that this might not be a problem. However, the autocorrelation is sufficiently

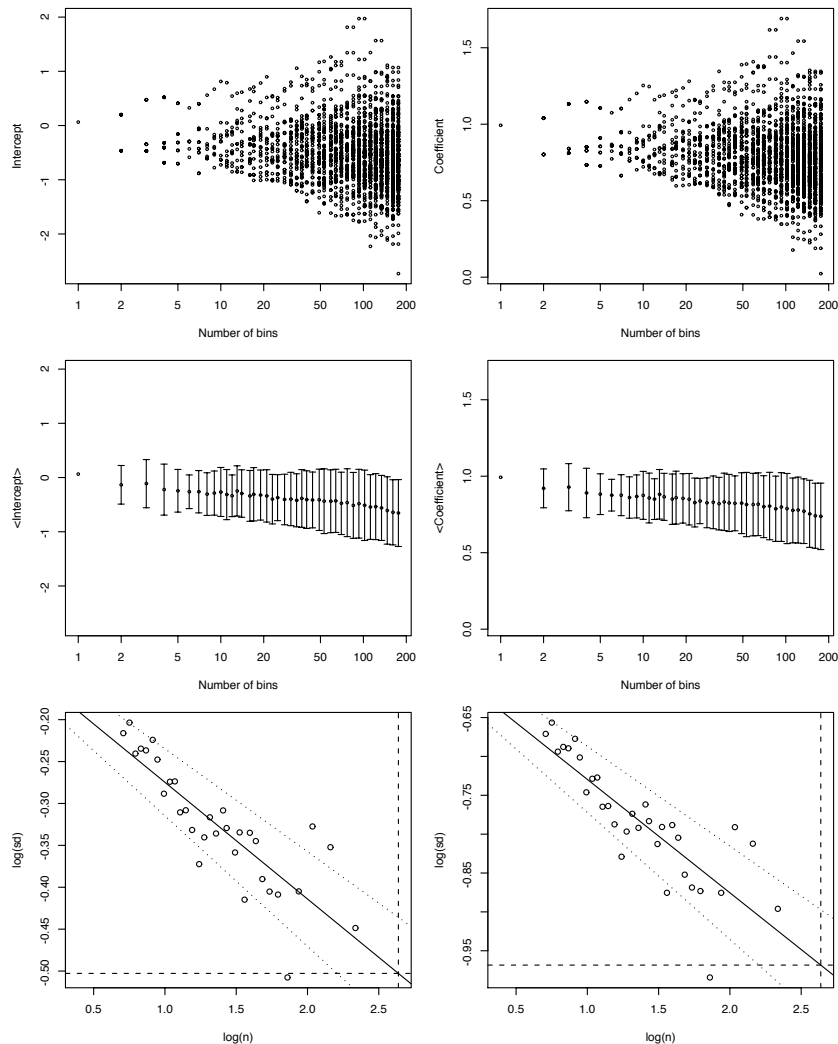


FIG. 3: Subsample analysis of regression of predicted vs. actual spread. To get a better feeling for the true errors in this estimation (as opposed to standard errors which are certainly too small), we divide the data into subsamples (using the same temporal period for each stock) and apply the regression to each subsample. (a) (top left) shows the results for the intercept, and (b) (top right) shows the results for the slope. In both cases we see that progressing from right to left, as the subsamples increase in size, the estimates become tighter. (c) and (d) (next row) shows the mean and standard deviation for the intercept and slope. We observe a systematic tendency for the mean to increase as the number of bins decreases. (e) and (f) show the logarithm of the standard deviations of the estimates against $\log n$, the number of each points in the subsample. The line is a regression based on binnings ranging from $m = N$ to $m = 10$ (lower values of m tend to produce unreliable standard deviations). The estimated error bar is obtained by extrapolating to $n = N$. To test the accuracy of the error bar, the dashed lines are one standard deviation variations on the regression, whose intercepts with the $n = N$ vertical line produce high and low estimates.

regression	estimated	standard	bootstrap	low	high
spread intercept	0.06	0.21	0.29	0.25	0.33
spread slope	0.99	0.08	0.10	0.09	0.11
diffusion intercept	2.43	1.22	1.76	1.57	1.97
diffusion slope	1.33	0.19	0.25	0.23	0.29

TABLE II: A summary of the bootstrap error analysis described in the text. The columns are (left to right) the estimated value of the parameter, the standard error from the cross sectional regression in Fig. 6, the one standard deviation error bar estimated by the bootstrapping method, and the one standard deviation low and high values for the extrapolation, as shown in Figs. 3(e-f) and 4(e-f).

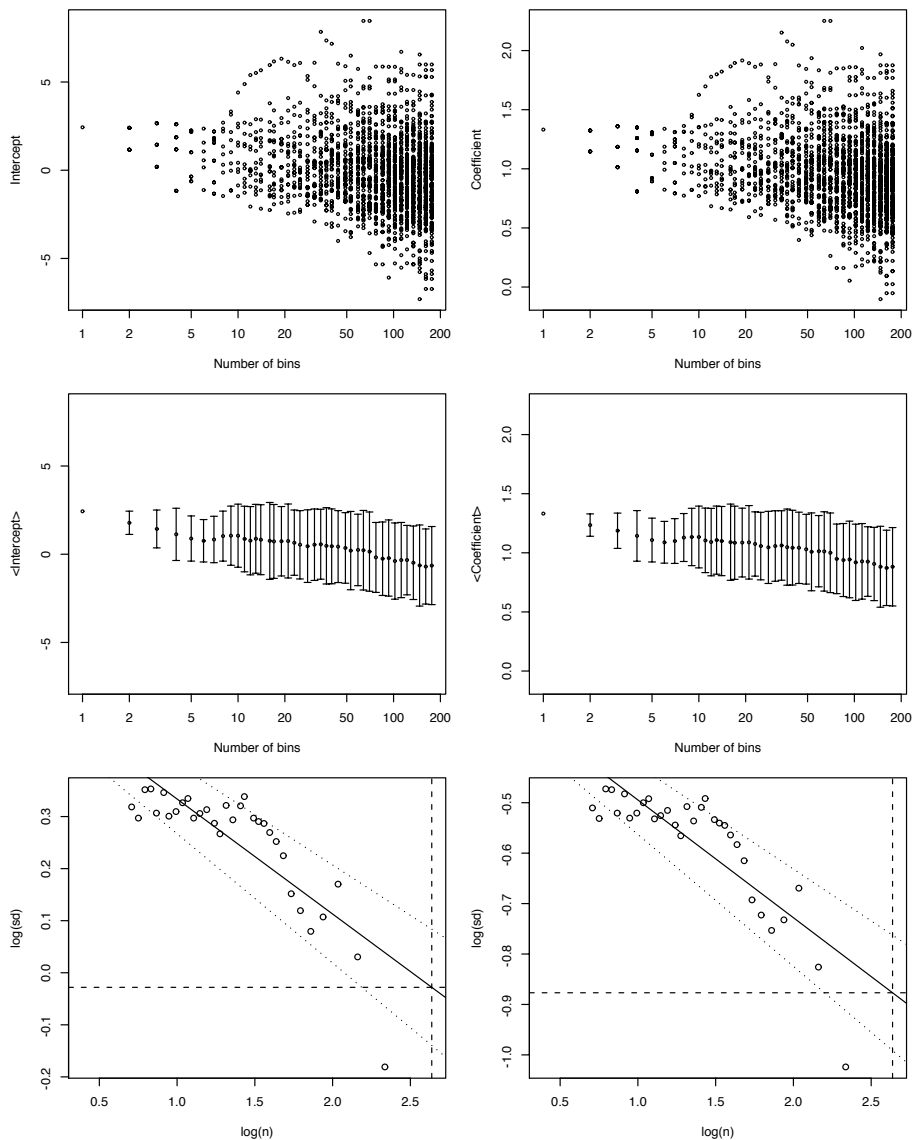


FIG. 4: Subsample analysis of regression of predicted vs. actual price diffusion (see Fig. 6), similar to the previous figure for the spread. The scaling of the errors is much less regular than it is for the spread, so the error bars are less accurate.

strong that its effect is still significant, particularly for smaller market impacts, and must be taken into account.

To cope with this we assign error bars to each bin using the variance plot method described in, for example, Beran [6], Section 4.4. This is a more straightforward version of the method discussed in Section (V). The sample of size $N = 434$ is divided into m subsamples of n points adjacent in time. We compute the mean for each subsample, vary n , and compute the standard deviation of the means across the $m = N/n$ subsamples. We then make use of theorem 2.2 from Beran [6] that states that the error in the n sample mean of a long-memory process is $\hat{e} = \sigma n^{-\gamma}$, where γ is a positive coefficient related to the Hurst exponent and σ is the standard deviation. By plotting the standard deviation of the m estimated inter-

cepts as a function of n we estimate γ and extrapolate to $n =$ sample length to get an estimate of the error in the full sample mean. An example of an error scaling plot for one of the bins of the market impact is given in Fig. 6.

A central question about Fig. 5 is whether the data for different stocks collapse onto a single curve, or whether there are statistically significant idiosyncratic variations from stock to stock. From the results presented in Fig. 5 this is not completely clear. Most of the stocks collapse onto the curve for the pooled data (or the pooled data set with themselves removed). There are a few that appear to make statistically significant variations, at least if we assume that the mean value of the bins for different order size levels are independent. However, they are most definitely *not* independent, and this non-independence

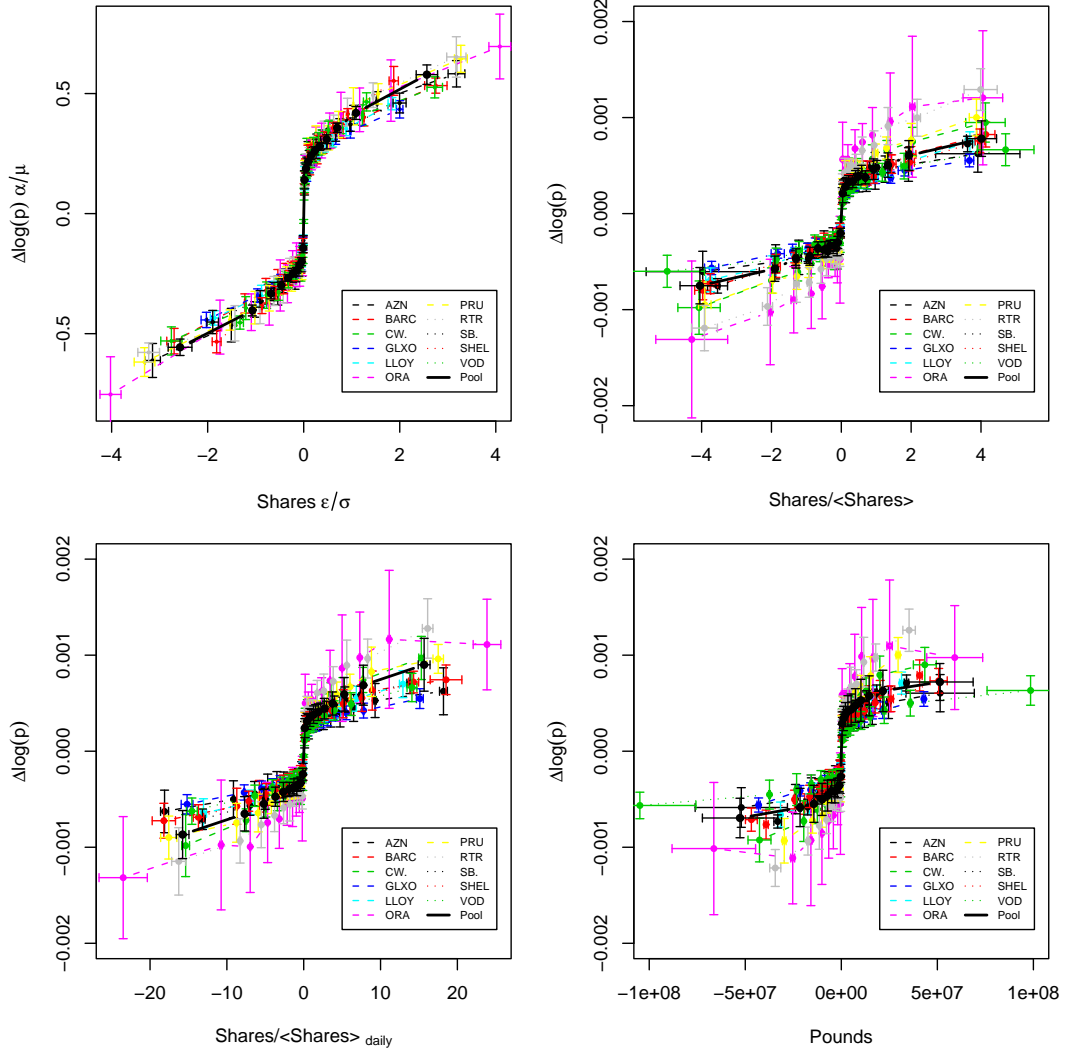


FIG. 5: Market impact collapse under 4 kinds of axis rescaling. In each case we plot a normalised version of the order size on the horizontal axis vs. a (possibly normalised) average market impact $\log(p_{t+1}) - \log(p_t)$ on the vertical axis. (a) (top left) collapse using non-dimensional units based on the model; (b) (top right) order size is normalised by its mean value for the sample. (c) (bottom right) order size is normalised the average daily volume. (d) (bottom right) Order size is multiplied by the current best midpoint price, making the horizontal axis the monetary value of the trade.

is difficult to model. In any case, the variations are always fairly small, not much larger than the error bars. Thus the collapse gives at least a good approximate understanding of the market impact, even if there are some small idiosyncratic variations it does not capture.

D. Market impact in log-log coordinates

If we fit a function of the form $\phi(\omega) = K\omega^\beta$ to the market impact curve, we get $\beta = 0.26 \pm 0.02$ for buy orders and $\beta = 0.23 \pm 0.02$ for sell orders, as shown in Fig. 7. The functional form of the market impact we observe here is not in agreement with a recent theory by

Gabaix et al. [33], which predicts $\beta = 0.5$. While the error bars given are standard errors, and are certainly too optimistic, it is nonetheless quite clear that the data are inconsistent with $\beta = 1/2$, as discussed in Ref. [29]. This relates to an interesting debate: The theory for average market impact put forth by Gabaix et al. follows traditional thinking in economics, and postulates that agents optimize their behavior to maximize profits, while the theory we test here assumes that they behave randomly, and that the form of the average market impact function is dictated by the statistical mechanics of price formation.

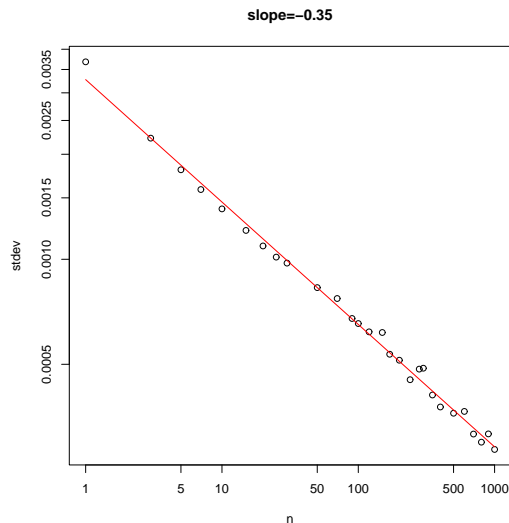


FIG. 6: The variance plot procedure used to determine error bars for mean market impact conditional on order size. The horizontal axis n denotes the number of points in the m different samples, and the vertical axis is the standard deviation of the m sample means. We estimate the error of the full sample mean by extrapolating n to the full sample length.

VII. EXTENDING THE MODEL

In the interest of full disclosure, and as a stimulus for future work, in this section we detail the ways in which the current model does not accurately match the data, and sketch possible improvements. This model was intended to describe a few average statistical properties of the market, some of which it describes surprisingly well. However, there are several aspects that it does not describe well. Fixing the problems with this models require a more sophisticated model of order flow, including a more realistic model of price dependence in order placement and cancellations [12, 65], long-memory properties [11, 47] and the relationship of the different components of the order flow to each other. This is a much harder problem, and is likely to require a more complicated model. Members of our group are actively working on this problem. While this will certainly have many advantages over the current model, it will also have the disadvantage of introducing more free parameters and thereby complicating the scaling laws (and making the possibility for analytic results more remote).

One of the major ways in which this model is not realistic concerns price diffusion. Real price increments are roughly white, i.e., they are roughly uncorrelated. One might naively think that under IID Poisson order flow, price increments should also be IID. However, due to the coupling of boundary conditions for the buy market order/sell limit order process to those of the sell market order/buy limit order process, this is not the case. Be-

cause of the fact that supply and demand tend to build (i.e. the depth of standing limit orders increases) as one moves away from the center of the book, price reversals are more common than price changes in the same direction. As a result, the price increments generated by this model are more anti-correlated than those of real price series. This has an interesting consequence: If we add the assumption of market efficiency, and assume that real price increments must be white, it implies that real order flow should be positively autocorrelated in order to compensate for the anticorrelations induced by the continuous double auction. This has indeed subsequently been observed to be the case [11, 47].

One of the side effects of this anticorrelation of prices is that it implies that there exist arbitrage opportunities that can be taken advantage of by an intelligent agent. A separate study of these arbitrage opportunities makes it clear that they are not risk-free in the sense usually used in economics. That is, taking advantage of them necessarily involves taking risk, and they do not permit arbitrarily large profits – returns decrease with the size of investment and eventually go to zero. Exploring the nature of these arbitrage opportunities, and the effect that exploiting them has on prices is one of the directions in which this model can be improved (one that is being actively explored). However, we do not feel that the existence of such arbitrage opportunities (which in our opinion mimic those of real markets) presents a serious problem for the purposes for which we are using this model.

In the following bullets we summarize the main directions in which members of our group are working to improve this model.

- *Price diffusion.* The variance of real prices obeys the relationship $\sigma^2(\tau) = D\tau^{2H}$ to a good approximation for all values of τ , where $\sigma^2(\tau)$ is the variance of price changes or returns computed on timescale τ . The Hurst exponent H is close to and typically a little greater than 0.5. In contrast, under Poisson order flow, as already discussed above, due to the dynamics of the double continuous auction price formation process, prices make a strongly anti-correlated random walk. This means that the function $\sigma^2(\tau)$ is nonlinear. Asymptotically $H = 0.5$, but for shorter times $H < 0.5$. Alternatively, one can characterize this in terms of a timescale-dependent diffusion rate $D(\tau)$, so that the variance of prices increases as $\sigma^2(\tau) = D(\tau)\tau$. Refs. [21, 60] showed that the limits $\tau \rightarrow 0$ and $\tau \rightarrow \infty$ obey well-defined scaling relationships in terms of the parameters of the model. In particular, $D(0) \sim \mu^2\delta/\alpha^2\epsilon^{-1/2}$, and $D(\infty) \sim \mu^2\delta/\alpha^2\epsilon^{1/2}$. Interestingly, and for reasons we do not fully understand, the prediction of the short term diffusion rate $D(0)$ does a good job of matching the real data, as we have shown here, while $D(\infty)$ does a much poorer job.

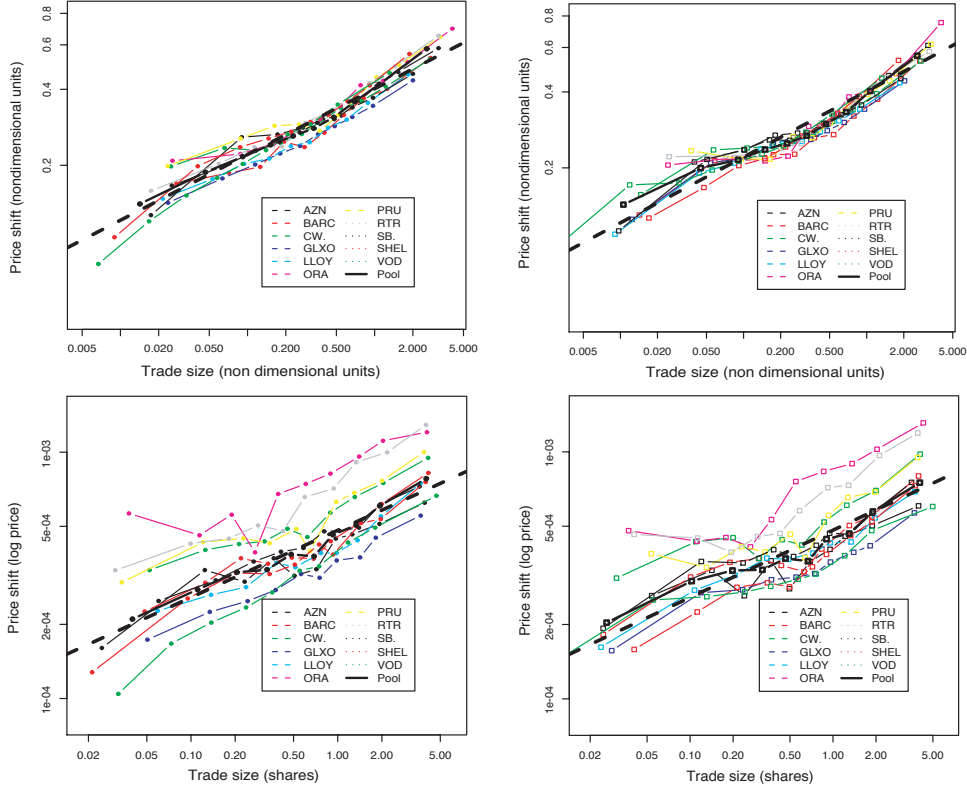


FIG. 7: The average market impact vs. order size plotted on log-log scale. The upper left and right panels show buy and sell orders in non-dimensional coordinates; the fitted line has slope $\beta = 0.26 \pm 0.02$ for buy orders and $\beta = 0.23 \pm 0.02$ for sell orders. In contrast, the lower panels show the same thing in dimensional units, using British pounds to measure order size. Though the exponents are similar, the scatter between different stocks is much greater.

- *Market efficiency.* The question of market efficiency is closely related to price diffusion. The anti-correlations mentioned above imply a market inefficiency. We are investigating the addition of “low-intelligence” agents to correct this problem.
- *Correlations in spread and price diffusion.* We have already discussed in Section (V) the problems that the autocorrelations in spread and price diffusion create for comparing the theory to the model on a daily scale. This is related to the fact that this model does not correctly capture either the fat tails of price fluctuations or the long-memory of volatility.
- *Lack of dependence on granularity parameter.* In Section (VIA) we discuss the fact that the model

predicts more variation with the granularity parameter than we observe. Apparently the Poisson-based non-dimensional coordinates work even better than one would expect. This suggests that there is some underlying simplicity in the real data that we have not fully captured in the model.

Although in this paper we are stressing the fact that we can make a useful theory out of zero-intelligence agents, we are certainly not trying to claim that intelligence doesn’t play an important role in what financial agents do. Indeed, one of the virtues of this model is that it provides a benchmark to separate properties that are driven by the statistical mechanics of the market institution from those that are driven by conditional intelligent behavior.

[1] Y. Amihud and H. Mendelson. Dealership market: Market-making with inventory. *Journal of Financial Economics*, 8(1):31–53, 1980.
 [2] T. G. Andersen. Return volatility and trading volume:

An information flow interpretation of stochastic volatility. *Journal of Finance*, 51(1):169–204, 1996.
 [3] J. J. Angel. Limit versus marker orders, 1994.
 [4] P. Bak, M. Paczuski, and M. Shubik. Price variations in

- a stock market with many agents, August 1996. Copies available from: Yale University, Cowles Foundation Library, Box 208281, New Haven, CT 06520.
- [5] G. Barenblatt. *Dimensional Analysis*. Gordon and Breach, New York, 1987.
 - [6] J. Beran. *Statistics for long-memory processes*. Chapman & Hall, New York, 1994.
 - [7] M. E. Blume, D. B. Keim, and S. A. Patel. Returns and volatility of low-grade bonds: 1977-1989. *Journal of Finance*, 46(1):49–74, 1991.
 - [8] T. Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31:307–327, 1986.
 - [9] T. Bollerslev, R. Y. Chou, and K. F. Kroner. Arch modeling in finance. *Journal of Econometrics*, 52(5-59), 1992.
 - [10] T. Bollerslev, I. Domowitz, and J. Wang. Order flow and the bid-ask spread: An empirical probability model of screen-based trading. *Journal of Economic Dynamics and Control*, 21(8-9):1471–1491, 1997.
 - [11] J.-P. Bouchaud, Y. Gefen, M. Potters, and M. Wyart. Fluctuations and response in financial markets: The subtle nature of “random” price changes. *Quantitative Finance*, 4(2):176–190, 2004.
 - [12] J.-P. Bouchaud, M. Mezard, and M. Potters. Statistical properties of the stock order books: empirical results and models. *Quantitative Finance*, 2(4):251–256, 2002.
 - [13] S. Chakravarty and C. W. Holden. An integrated model of market and limit orders. *Journal of Financial Intermediation*, 4(3):213–41, 1995.
 - [14] D. Challet and R. Stinchcombe. Analyzing and modeling 1+1d markets. *Physica A*, 300(1-2):285–299, 2001.
 - [15] C. Chiarella and G. Iori. A simulation analysis of the microstructure of double auction markets. *Quantitative Finance*, 2:346–353, 2002.
 - [16] W. S. Choi, S. B. Lee, and P. I. Yu. Estimating the permanent and transitory components of the bid/ask spread. In C.-F. e. Lee, editor, *Advances in investment analysis and portfolio management. Volume 5*. 1998.
 - [17] P. K. Clark. Subordinated stochastic process model with finite variance for speculative prices. *Econometrica*, 41(1):135–155, 1973.
 - [18] K. J. Cohen, R. M. Conroy, and S. F. Maier. Order flow and the quality of the market. In Y. Amihud, T. Ho, and R. A. Schwartz, editors, *Market Making and the Changing Structure of the Securities Industry*. Rowman & Littlefield, Lanham, 1985.
 - [19] K. J. Cohen, S. F. Maier, R. A. Schwartz, and D. K. Whitcomb. Transaction costs, order placement strategy, and existence of the bid-ask spread. *Journal of Political Economy*, 89(2):287–305, 1981.
 - [20] T. E. Copeland and D. Galai. Information effects on the bid-ask spread. *Journal of Finance*, 38(5):1457–69, 1983.
 - [21] M. G. Daniels, J. D. Farmer, G. Iori, and E. Smith. Quantitative model of price diffusion and market friction based on trading as a mechanistic random process. *Physical Review Letters*, 10(90), 2003.
 - [22] H. Demsetz. The cost of transacting. *The Quarterly Journal of Economics*, 82:33–53, 1968.
 - [23] I. Domowitz and J. Wang. Auctions as algorithms: Computerized trade execution and price discovery. *Journal of Economic Dynamics and Control*, 18(1):29–60, 1994.
 - [24] D. Easley and M. O’Hara. Price, trade size, and information in securities markets. *Journal of Financial Economics*, 19(1):69–90, 1987.
 - [25] D. Easley, M. O’Hara, and G. Saar. How stock splits affect trading: A microstructure approach. *Journal of Financial and Quantitative Analysis*, 36(1):25–51, 2001.
 - [26] D. Eliezer and I. I. Kogan. Scaling laws for the market microstructure of the interdealer broker markets, 1998.
 - [27] R. F. Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of united-kingdom inflation. *Econometrica*, 50(4):987–1007, 1982.
 - [28] T. W. Epps and M. L. Epps. The stochastic dependence of security price changes and transaction volumes: Implications for the mixture-of-distributions hypothesis. *Econometrica*, 44:305–321, 1976.
 - [29] J. D. Farmer, P. Patelli, and I. Zovko. The predictive power of zero intelligence. Technical report, Santa Fe Institute, 2003.
 - [30] J. Fleming, C. Kirby, and B. Ostdiek. The economic value of volatility timing. *Journal of Finance*, 56(1):329–52, 2001.
 - [31] T. Foucault. Order flow composition and trading costs in a dynamic limit order market. *Journal of Financial Markets*, 2(2):99–134, 1999.
 - [32] T. Foucault, O. Kadan, and E. Kandel. Limit order book as a market for liquidity, 2001.
 - [33] X. Gabaix, P. Gopikrishnan, V. Plerou, and H. E. Stanley. A theory of power-law distributions in financial market fluctuations. *Nature*, 423:267–270, 2003.
 - [34] A. R. Gallant, P. E. Rossi, and G. E. Tauchen. Nonlinear dynamic structures. *Econometrica*, 61:871–908, 1993.
 - [35] R. A. Gallant, P. E. Rossi, and G. E. Tauchen. Stock prices and volumes. *Review of Financial Studies*, 5:199–242, 1992.
 - [36] T. J. George, G. Kaul, and M. Nimalendran. Estimation of the bid-ask spread and its components: A new approach. *Review of Financial Studies*, 4(4):623–56, 1991.
 - [37] L. L. Glosten. Estimating the components of the bid/ask spread. *Journal of Financial Economics*, 21:123–142, 1988.
 - [38] L. R. Glosten. Competition and the set of allowable prices, October 1995. Copies available from: Graduate School of Business, 6N Uris Hall, Columbia University, 3022 Broadway, New York, NY 10027. Website: www.columbia.edu/dlc/wp/.
 - [39] L. R. Glosten and P. R. Milgrom. Bid, ask, and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1):71–100, 1985.
 - [40] L. Harris. Optimal dynamic order submission strategies in some stylized trading problems. *Financial Markets, Institutions and Instruments*, 7(2):1–75, 1998.
 - [41] J. Hasbrouck. Trades, quotes, inventories, and information. *Journal of Financial Economics*, 22(2):229–52, 1988.
 - [42] J. Hasbrouck. Intraday price formation in u.s. equity index markets. *Journal of Finance*, 58(6):2375–99, 2003.
 - [43] T. Ho and H. R. Stoll. Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics*, 9(1):47–73, 1981.
 - [44] R. D. Huang and H. R. Stoll. The components of the bid-ask spread: A general approach. *Review of Financial Studies*, 10(4):995–1034, 1997.
 - [45] H. Jang and P. C. Venkatesh. Consistency between predicted and actual bid-ask quote-revisions. *Journal of Finance*, 46(1):433–46, 1991.
 - [46] C. G. Lamoureux and W. D. Lastrapes. Endogenous

- trading volume and momentum in stock-return volatility. *Journal of Business and Economic Statistics*, 12:253–260, 1994.
- [47] F. Lillo and J. D. Farmer. The long memory of the efficient market. Working paper 03-12-066, Santa Fe Institute, 2003.
- [48] A. W. Lo. Bubble, rubble, finance in trouble? *The Journal of Psychology and Financial Markets*, 3(2):76–86, 2002.
- [49] LSEbulletin. Sets four years on. Technical report, London Stock Exchange, 2001.
- [50] A. Madhavan, M. Richardson, and M. Roomans. Why do security prices change? a transaction-level analysis of nyse stocks. *The Review of Financial Studies*, 10(4):1035–1064, 1997.
- [51] S. Maslov. Simple model of a limit order-driven market. *Physica A-Statistical Mechanics and Its Applications*, 278(3-4):571–578, 2000.
- [52] H. Mendelson. Market behavior in a clearing house. *Econometrica*, 50(6):1505–24, 1982.
- [53] C. A. Parlour. Price dynamics in limit order markets. *Review of Financial Studies*, 11(4):789–816, 1998.
- [54] C. A. Parlour and D. J. Seppi. Liquidity-based competition for order flow. *Review of Financial Studies*, 16(2):301–43, 2003. ja href="http://rfs.oupjournals.org";Publisher's URL;aj.
- [55] P. C. Reiss and I. M. Werner. Transaction costs in dealer markets: Evidence from the london stock exchange, May 1994. Copies available from: National Bureau of Economic Research, 1050 Massachusetts Avenue, Cambridge, MA 02138.
- [56] K. Rock. The specialist's order book and price anomalies, 1990.
- [57] R. Roll. A simple implicit measure of the effective bid-ask spread in an efficient market. *Journal of Finance*, 39(4):1127–39, 1984.
- [58] D. J. Seppi. Liquidity provision with limit orders and a strategic specialist. *Review of Financial Studies*, 10(1):103–50, 1997.
- [59] F. Slanina. Mean-field approximation for a limit order driven market model. *Physical Review E*, 6405(5), 2001. Part 2.
- [60] E. Smith, J. D. Farmer, L. Gillemot, and S. Krishnamurthy. Statistical theory of the continuous double auction. *Quantitative Finance*, 3(6):481–514, 2003.
- [61] H. R. Stoll. The supply of dealer services in securities markets. *Journal of Finance*, 33(4):1133–51, 1978.
- [62] L. H. Tang and G. S. Tian. Reaction-diffusion-branching models of stock price fluctuations. *Physica A*, 264(3-4):543–550, 1999.
- [63] G. E. Tauchen and M. Pitts. The price variability-volume relationship on speculative markets. *Econometrica*, 51:485–505, 1983.
- [64] S. M. Tinic. The economics of liquidity services. *Quarterly Journal of Economics*, 86(1):79–93, 1972.
- [65] I. Zovko and J. D. Farmer. The power of patience; a behavioral regularity in limit order placement. *Quantitative Finance*, 2(5):387–392, 2002.