The Predictive Power of Zero Intelligence in Financial Markets

J. Doyne Farmer[1], Paolo Patelli[1,2], 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

Original version Sept. 9, 2003; this version Nov. 16, 2003

Abstract

Standard models in economics are based on intelligent agents that maximize utility. However, there may be situations where constraints imposed by market institutions are more important than intelligent agent behavior. We use data from the London Stock Exchange to test a simple model in which zero intelligence agents place orders to trade at random. The model treats the statistical mechanics of the interaction of order placement, price formation, and the accumulation of stored supply and demand, and makes predictions that can be stated as simple expressions in terms of measurable quantities such as order arrival rates. The agreement between model and theory is excellent, explaining 96% of the variance of the bid-ask spread across stocks and 76% of the price diffusion rate. We also study the market impact function, describing the response of prices to orders. The non-dimensional coordinates dictated by the model collapse data from different stocks onto a single curve, suggesting a corresponding understanding of supply and demand. Thus, it appears that the price formation mechanism strongly constrains the statistical properties of the market, playing a more important role than the strategic behavior of agents.

Contents

1 Introduction

Since the nineteenth century one of the classic questions in economics has been, "What determines supply and demand?". Similarly, since Bachelier [1] introduced the random walk model for prices in 1900, a question that has been crying out for an answer has been, "What determines the diffusion rate of prices?". Standard models in economics, which are based on rational utility maximizing agents, have had only limited success in addressing these questions. In this paper we demonstrate that a model built on the opposite approach – that agents are of zero intelligence, and simply place orders to trade at random – can successfully address these questions and others, providing one properly models the statistical mechanics of price formation.

Traditionally economics has devoted considerable effort to modeling the strategic behavior and expectations of agents. While no one would dispute that this is important, it has also been pointed out that some aspects of economics are independent of the agent model. For example, Becker [2] showed that a budget constraint is sufficient to guarantee the proper slope of supply and demand curves, and Gode and Sunder [3] demonstrated that if one replaces the students in a standard classroom economics experiment by zero-intelligence agents, the zero-intelligence agents perform surprisingly well. In this paper we show that this principle can be dramatically more powerful, and can make surprisingly accurate quantitative predictions, if one models the statistical mechanics of the price formation process. In particular, we test a zero-intelligence statistical mechanics model due to Daniels et al. [4, 5], which builds on earlier work in financial economics [6, 7, 8, 9] and physics [10, 11, 12, 13, 14] (See also interesting subseqent work [15, 16]). This model added to the prior literature by constructing and approximately solving a simple model for price setting that makes quantitative, testable predictions about fundamental market properties, many of which can be expressed as simple algebraic formulas.

1.1 Continuous double auction

The model of Daniels et al. [4] assumes a continuous double auction, which is the most widely used method of price formation in modern financial markets [5]. There are two fundamental kinds of trading orders: Impatient traders submit *market orders*, which are requests to buy or sell a desired number of shares immediately at the best available price. More patient traders submit *limit orders*, which include the worst allowable price for the transaction. Limit orders may fail to result in an immediate transaction, in which case they are stored in a queue called the *limit order book*, illustrated in Fig. 1. As each buy order arrives it is transacted against accumulated sell limit orders that have a lower selling price, in priority of price and arrival time. Similarly for sell orders. The lowest selling price offered in the book at any point in time is called the *best ask*, $a(t)$, and the highest buying price the *best bid*, $b(t)$.

1.2 Review of model

The model assumes that two types of zero intelligence agents place and cancel orders randomly (see Fig. 1). Impatient agents place market orders of size σ, at a rate μ *shares per time*. Patient agents place limit orders of the same size σ, with a constant rate density α *shares per price per time*, and queued limit orders are canceled at a constant rate δ, with dimensions of *1/time*. Prices change in discrete increments called *ticks*, of size dp. To keep the model as simple as possible, the rates are equal for buying and selling, and order placement and cancellation are Poisson processes. All of these processes are independent except for coupling through their boundary conditions: Buy limit orders arrive with a constant density α over the semi-infinite interval $-\infty < p < a(t)$, where p is the logarithm of the price, and sell limit orders arrive with constant density α on the semi-infinite interval $b(t) < p < \infty$. As new orders arrive they may alter the best prices $a(t)$ and $b(t)$, which in turn changes the boundary conditions for subsequent limit order placement. For example, the arrival of a buy limit order inside the spread will alter the best bid $b(t)$, which immediately alters the boundary condition for subsequent sell limit order placement. As a result of the random order arrival processes, $a(t)$ and $b(t)$ each make random walks, but because of coupling of the buying and selling processes the bid-ask *spread* $s(t) \equiv a(t) - b(t)$ is a stationary random variable. It is this feedback between order placement and price diffusion that makes this model interesting, and de-

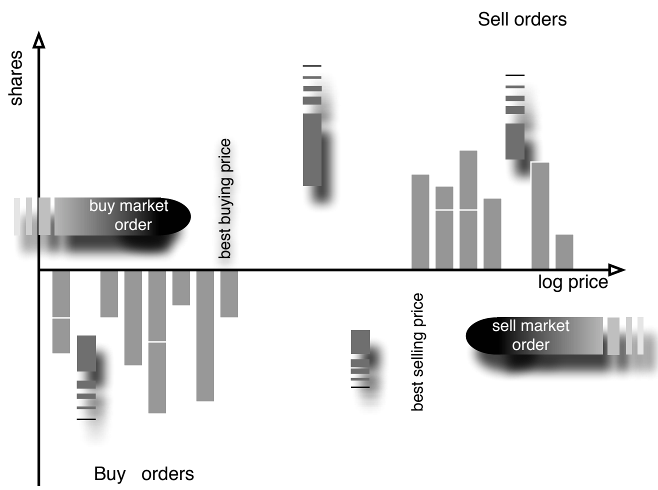

Figure 1: A random process model of the continuous double auction. Stored limit orders are shown stacked along the price axis, with sell orders (supply) stacked above the axis at higher prices and buy orders (demand) stacked below the axis at lower prices. New sell limit orders are visualized as randomly falling down, and new buy orders as randomly "falling up". New sell orders can be placed anywhere above the best buying price, and new buy orders anywhere below the best selling price. Limit orders can be removed spontaneously (e.g. because the agent changes her mind or the order expires) or they can be removed by market orders of the opposite type. This can result in changes in the best prices, which in turn alters the boundaries of the order placement process. It is this feedback between order placement and price formation that makes this model interesting, and its predictions non-trivial.

spite its apparent simplicity, quite difficult to understand analytically [5].

While $a(t)$ and $b(t)$ make random walks, the increments of their random walks are strongly anticorrelated. This is a good example of how the properties of this model are not simple to understand. 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. Because of the fact that supply and demand tend to build 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 strongly anti-correlated, much more so than real price series. This has interesting consequence: If we add the assumption of strong market efficiency, and assume that real price increments must be white, it implies that real order flow should be strongly positively correlated in order to compensate for the anticorrelations induced by the continuous double auction. This has indeed been observed to be the case [21, 22].

This model has been analyzed using both simulation and two different mean field theories [5]. The analytic prediction of the spread based on mean field theory match the simulations quite well. They also do a fairly good job of matching the simulation results for the average supply and demand curves defined by the limit order book. The analytic methods are not as successful in predicting the price diffusion rate, and so here we will fall back on simulation results.

One of the virtues of this model is that we can make approximate predictions of several of its properties with almost no work using dimensional analysis. This also greatly simplifies the analysis and understanding of the model, which as we will see is particularly useful for understanding market impact. There are three fundamental dimensional quantities describing everything in this model: *shares*, *price*, and *time*. This is apparent already from Fig. 1, whose axes have units of price and shares – time is the third dimension. On the other hand, there are five parameters defined in the model. When the dimensional constraints between the parameters are taken into account, this leaves only two independent degrees of freedom in the model.

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. In particular, 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 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, and as we will see, many quantities are much more well-behaved and easily understood when plotted in non-dimensional coordinates than they are otherwise.

The remaining two 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 *delta* that keep the two non-dimensional parameters constant will result in exactly the same market properties. One of the interesting properties 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 [5].

Another one of the virtues of this model is that it results in simple scaling laws relating the parameters of the model to fundamental properties such as the average bid-ask spread, the price diffusion rate, and average supply and demand curves. Rather than reviewing these here, we will review each of them as we compare to the data. However, we would like to emphasize that the construction of the model and all the predictions derived from the model were made before looking at the data.

2 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 electronic 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 while 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 [19]. Our analysis involves only the data from the electronic order book. We chose 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 the 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 1.

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;

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 1: 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.

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 stock splits and tick price changes during the 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.

3 Measurement of model parameters

Our goal is to compare the predictions of the model with real data. The parameters of the model are stated in terms of order arrival rates, cancellation rate, order size, and tick size. We chose an appropriate time interval and measure the parameters over that interval, and then compare to the properties of the market over that same interval.

Reconstructing the limit order book on a moment-by-moment basis makes it clear that the properties of the market tend to be relatively stationary during each day, changing more dramatically at the beginning and at the end of day. It is therefore natural to measure each parameter for each stock on each day. 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 a effective market order, while the non-transacted part (if any) is counted as a 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 affect on prices, and in any case, make up only a very small fraction of the order flow, typically less than 1%. Note that throughout the remainder of the paper, we will tend to 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 price intervals are assumed to be logarithmic in order to assure prices are always positive.

We measure the average value of the five parameters of the model, μ, α, δ, σ, and dp for each day, making the assumption that the parameters of the model are stationary within each day, but change from day to day. μ, σ, and dp are straightforward to measure, but there are problems in measuring α and δ that must be understood in order to properly interpret our results.

The parameter μ_t, which characterizes the average market order arrival rate on day t, is straightforward to measure. It 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. Similarly, σ_t is the average limit order size[1] in shares for that day, and dp_t is just the tick size, which is fixed for each day but varies from stock to stock and for a given stock changes occasionally during the sample.

Measuring the cancellation rate δ_t and the limit order rate density α_t is more complicated, due to the highly simplified assumptions we have 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 at the best prices), with a density that falls off as a power law as a function of the distance Δ from the best prices [15, 20]. 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 as described 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 $\Delta = a - p$, where p is the logarithm of the limit price and a is the logarithm of the best selling price. Similarly for sell orders, $\Delta = p - b$, where b is the best buying price. (By using this convention we include all effective limit orders and guarantee that Δ is always positive). 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. 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 on day t. These choices are made in a compromise to include as much data as possible 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.

[1] The model assumes that the average size of limit orders and market orders is the same. For the real data this is not strictly true, though as seen in Table 1, it is a good approximation to within about 20%. For the purposes of the analysis we use the limit order size as the measure because for theoretical reasons we think this is more important than the market order size, but because the two are approximately the same, this will not make a significant difference in the results.

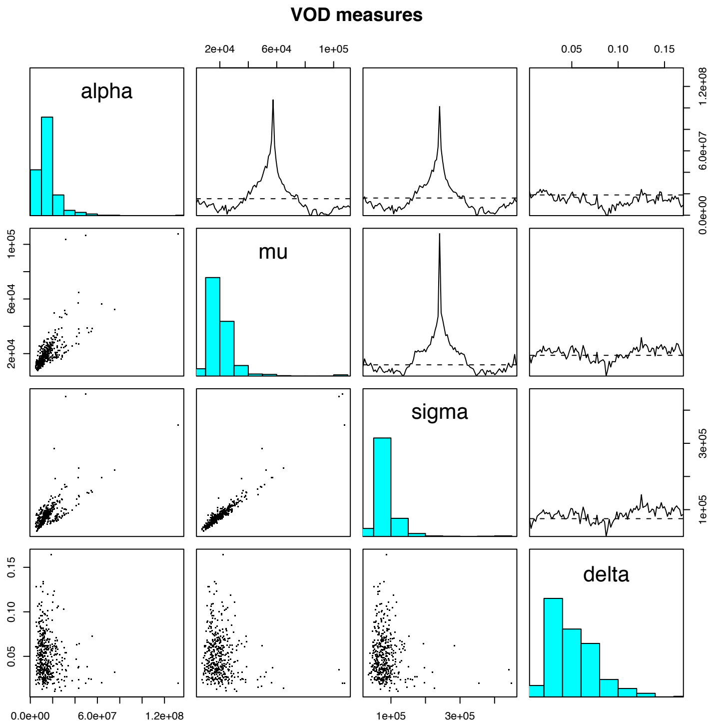

Figure 2: Density estimations and cross correlations for Vodafone between the four model parameter measures. On the diagonal we present the histogram of the corresponding parameter. Upper off-diagonal plots are the time cross correlation. We see that δ is uncorrelated with other measures, while the other three are quite correlated although without any noticeable lead-lag effects. The lower off-diagonal plots are scatter plots between the parameters. μ and α are particularly strongly correlated; fortunately, for the prediction of the spread their ratio is the most important quantity, and this correlation largely cancels out.

Similarly, to cope with the fact that in reality the average cancellation rate δ decreases [15] 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. Some simple diagnostics of the parameter estimates are presented in Fig. 2.

The above is obviously a crude simplification that allows us to use a model constructed for simplicity and ease of analysis. The arbitrary choices involved in choosing price intervals will introduce some uncertainty into the scale of the parameters, as discussed later when we present the results. We are also developing more realistic order placement and cancellation models. However, this model has the important advantage of simplicity, and in particular, we will see that its non-dimensional coordinates have unanticipated power.

4 Spread and price diffusion rate

4.1 Predicted vs. actual spread

The bid-ask spread is of central interest in financial markets because it is an important component of transaction costs. The mean value of the spread, predicted based on a mean field theory analysis of the model is $\hat{s} = (\mu/\alpha) f(\epsilon, dp/p_c)$, where f is a relatively slowly varying, monotonically increasing non-dimensional function that can be approximated as $f(\epsilon) = 0.28 + 1.11 * \epsilon^{3/4}$, neglecting the dp/p_c dependence [5]. To test this relationship, we measure the actual average spread $\bar{s}$ across the full time period for each stock, and compare to the predicted average spread $\hat{s}$ based on order flows. Spread is measured as the daily average of $\log(\text{bid}) - \log(\text{ask})$. The spread is measured after each event, with each event given equal weight. The opening auction is excluded. For reasons that are detailed in Section 4.4 we choose to test the model cross-sectionally.

To test our hypothesis that the predicted and actual values coincide, we perform a regression of the form $\log \bar{s} = A \log \hat{s} + B$. We use the free parameters A and B for hypothesis testing; apriori we expect $A = 1$ and $B = 0$. However, because of the parameter measurement problems described above, we do not expect the value of the intercept to be very meaningful, i.e. the hypothesis tested is that first of all this relationship is linear, and secondly that $A = 1$. (We used logarithms because the spread is positive and the log of the spread is approximately normally distributed).

The least squares regression, shown together with the data comparing the predictions to the actual values in Fig. 3, gives $A = 0.99 \pm 0.10$ and $B = 0.06 \pm 0.29$. We thus very strongly reject the null hypothesis that $A = 0$, indicating that the predictions are far better than random. More importantly, we are unable to reject the null hypotheses that $A = 1$. (In fact, we are also unable to reject $B = 0$, but this may be largely a matter of luck in our choice of Q_t^{upper}.) The regression has an $R^2 = 0.96$, so the adjusted model explains most of the variance. Note that because of long–memory effects and cross–correlations between stocks the error bars are much larger than they would be for IID, normally distributed data (see the discussion in Section 4.3).

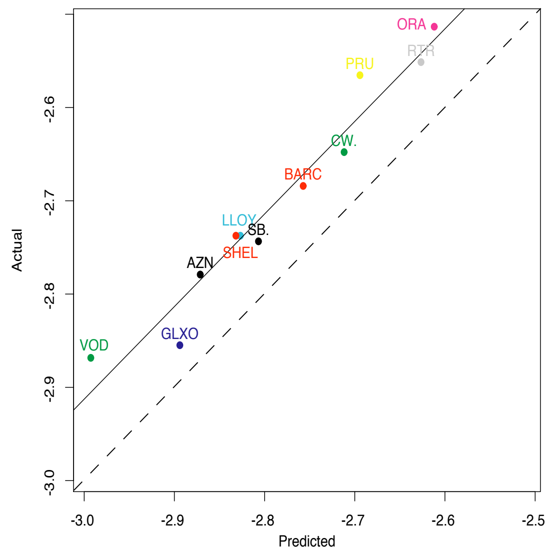

Figure 3: Regressions of predicted values based on order flow parameters vs. actual values for the log spread. The dots show the average predicted and actual value for each stock averaged over the full 21 month time period. The solid line is a regression; the dashed line is the diagonal, representing the model's prediction without any adjustment of slope or intercept.

4.2 Price diffusion rate

Another quantity of primary interest is the price diffusion rate, which drives the volatility of prices and is the primary determinant of financial risk. If we assume that prices make a random walk, then the diffusion rate measures the size and frequency of its increments. The variance V of an uncorrellated normal random walk after time t grows as $V(t) = Dt$, where D is the diffusion rate. We choose to measure the price diffusion rate rather than the volatility because it is a stationary quantity that provides a more fundamental description of the volatility process. This is the main free parameter in the Bachelier model, and while its value is essential for risk estimation and derivative pricing there is very little understanding of what determines it.

The measurement of the price diffusion rate requires some discussion. The variance of mid-point returns at time scale τ is defined as $\sigma^2(\tau) = \langle (m(t+$

$\tau)-m(t))^2\rangle$, where $\langle\cdot\rangle$ indicates an (event weighted) time average. For a random walk with stationary increments the variance increases with τ as

$$\sigma^2 = D\tau^{2H}, \tag{1}$$

where D is the diffusion rate. H is the Hurst exponent, which for an IID Gaussian random walk is $H = 1/2$. By measuring this variance at different intra-day time scales τ we estimate D using expression (1) by a linear regression weighted by the square root of the number of independent observations, and assuming $H = 1/2$. An example of this procedure is given in Fig. 4.

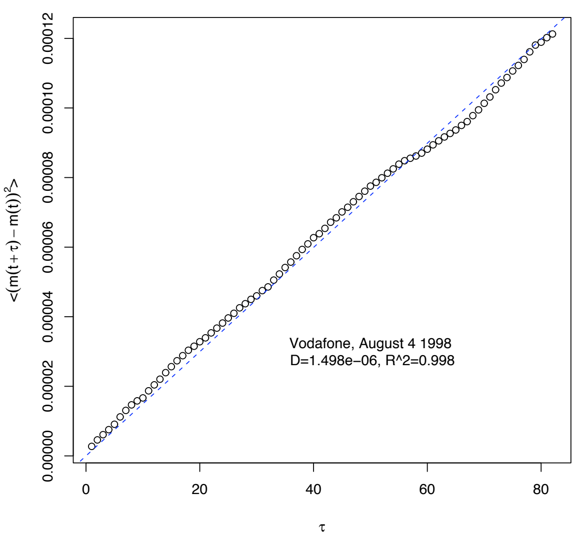

Figure 4: Illustration of the procedure for measuring the price diffusion rate for Vodafone (VOD) on August 4th, 1998. On the x axis we plot the time τ in units of ticks, and on the y axis the variance of mid-price diffusion $\sigma^2(\tau)$. According to the hypothesis that mid-price diffusion is an uncorrelated Gaussian random walk, the plot should be linear with slope equal to the diffusion constant (see equation (1)). 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.

The estimated slope of the variance plot is the diffusion rate D_t for day t. One must bear in mind that the price diffusion rate from day to day has substantial correlations, as illustrated in Fig. 5.

Numerical experiments indicate that the short term price diffusion rate predicted by the model is

$$\hat{D} = k\mu^{5/2}\delta^{1/2}\sigma^{-1/2}\alpha^{-2}, \tag{2}$$

where k is a constant. As for the spread, we compare this to the actual price diffusion rate $\bar{D}_i$ for each stock averaged over the 21 month period, and regress the logarithm of the predicted vs. actual values, as shown in Fig. 6.

The regression gives $A = 1.33 \pm 0.25$ and $B = 2.43 \pm 1.75$. Thus, we again strongly reject the null hypothesis that $A = 0$. We are still unable to reject the null hypothesis that $A = 1$ with 95% confidence, though there is some suggestion that the scaling of the model and the actual values are not quite the same. (This could happen if, for example, the scaling exponent predicted by the model of one or more of the order flow rates is too low). The reader should bear in mind that because of the long-memory effects and the cross-correlations between stocks, as discussed in Section 4.3, these error bars are difficult to determine. Although the results are not as good as for the spread, $R^2 = 0.76$, so the model still explains most of the variance.

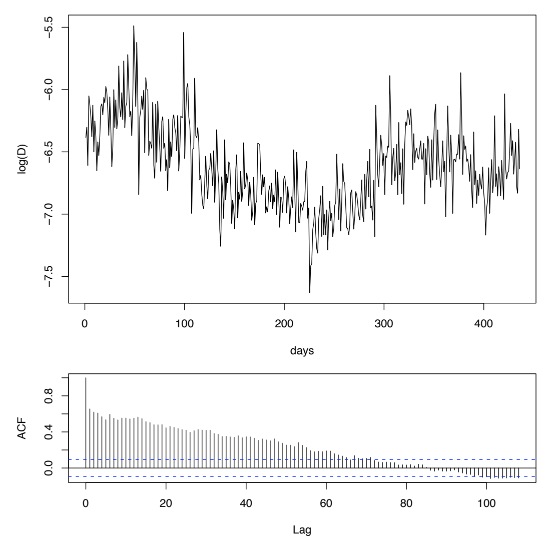

Figure 5: Time series (top) and autocorrelation function (bottom) for daily price diffusion rate D_t for Vodafone. Because of long-memory effects and the short length of the series, the long-lag coefficients are poorly determined; the figure is just to demonstrate that the correlations are quite large.

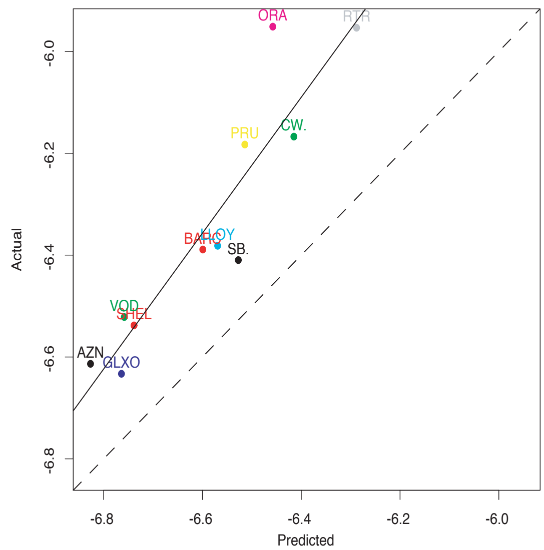

Figure 6: Regressions of predicted values based on order flow parameters vs. actual values for the logarithm of the price diffusion rate. The dots show the average predicted and actual value for each stock averaged over the full 21 month time period. The solid line is a regression; the dashed line is the diagonal, representing the model's prediction without any adjustment of slope or intercept.

4.3 Estimating the errors for the regressions

The error bars presented in the text are based on a bootstrapping method. We are driven to use this method because order flow variables, spread, and price diffusion rates all have slowly decaying positive autocorrelation functions. Indeed, it has recently been shown that order sign, order volume and liquidity as reflected by volume at the best price, are long-memory processes [21, 22]. In addition the spread, price diffusion rates, and parameters are highly cross-correlated between stocks. These effects complicate the statistical analysis, and make the assignment of error bars difficult.

The method we use is inspired by the variance plot method described in Beran [23], 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, $i = 1, \ldots, m$. We then vary m and study the scaling as shown in Figs. 7 and 8.

Figs. 7(a) and (b) illustrate this procedure for the spread, and Figs. 8(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 data with no auto- or cross-correlation we expect a line with slope $\gamma = -1/2$; instead we observe $\gamma > -1/2$. For example for the spread $\gamma \approx -0.19$. The smaller γ is an indication that this is a long memory process; see the discussion in Section (5.3).

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. 7 and 8. The results are summarized in Table 2.

One of the effects that is evident in Figs. 7(c-d) and 8(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 (4.4).

4.4 Longitudinal vs. cross-sectional tests

It is possible to test this model either longitudinally (across different time intervals for a given stock) or cross-sectionally (across different stocks over the same time period). We have applied tests of both types, but due to the very strong autocorrelations

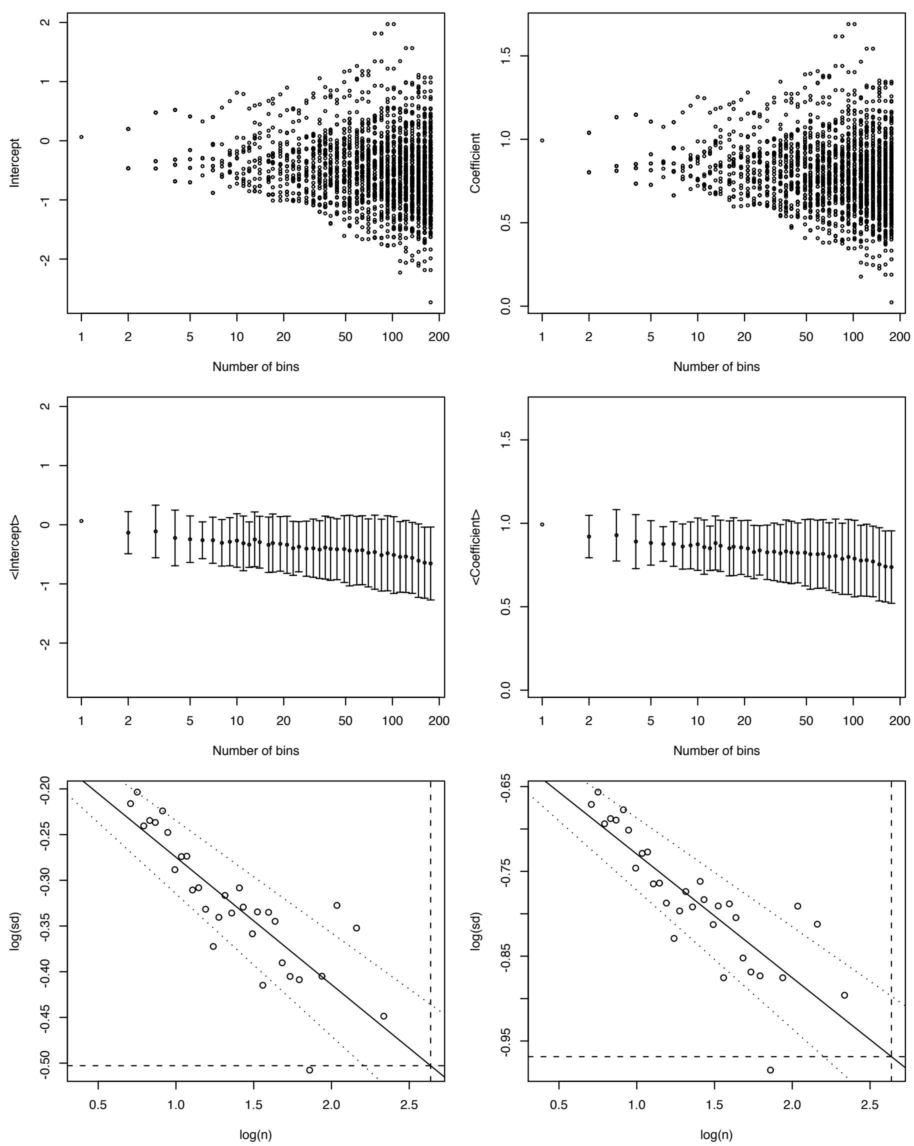

Figure 7: Subsample analysis of regression of predicted vs. actual spread. To get a better feeling for the true errors in this estimation (as opposed to standard errors which are certainly too small), we divide the data into subsamples (using the same temporal period for each stock) and apply the regression to each subsample. (a) (top left) shows the results for the intercept, and (b) (top right) shows the results for the slope. Each point is the value of the estimate for one of the bins. 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.

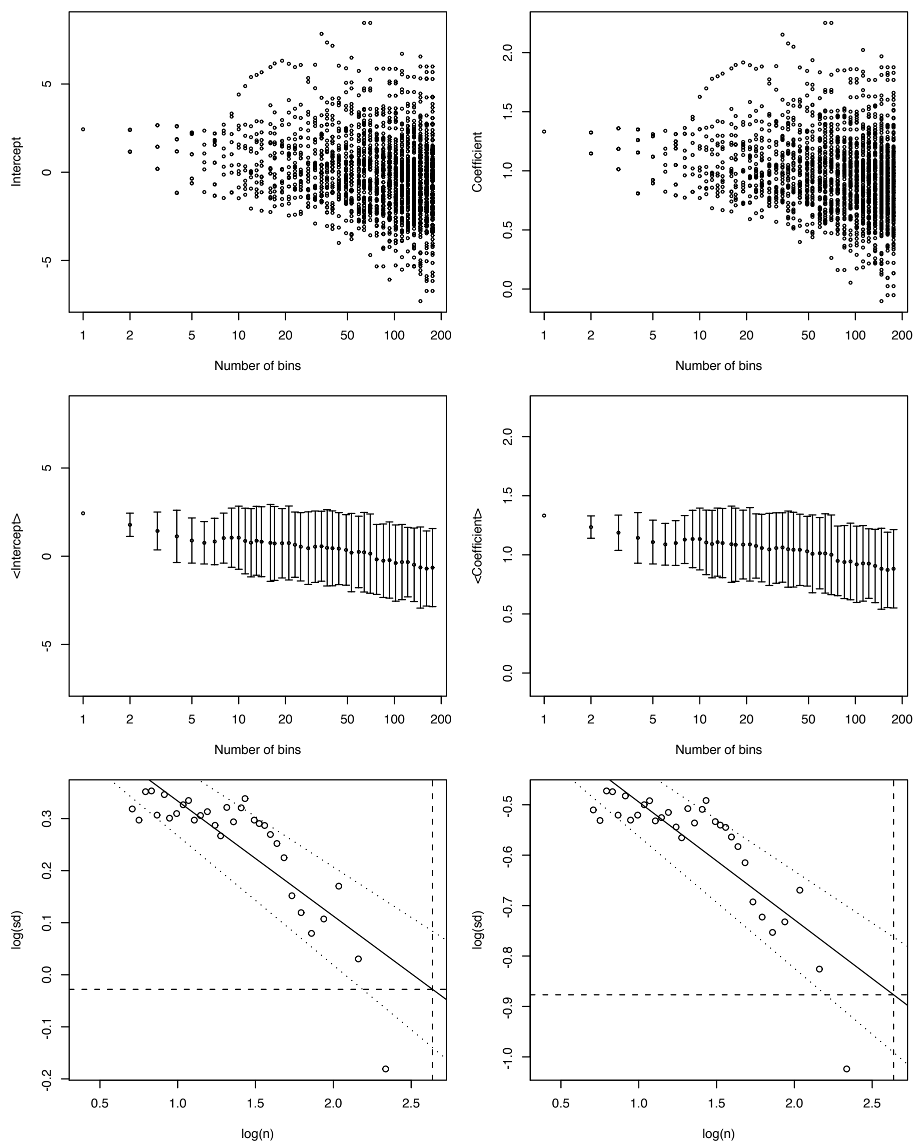

Figure 8: Subsample analysis of regression of predicted vs. actual price diffusion (see Fig. 6), similar to the previous figure for the spread. The scaling of the errors is much less regular than it is for the spread, so the error bars are less accurate.

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 2: 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. 7(e-f) and 8(e-f).

of the order flow rates, spread, and price diffusion rates, there are difficulties in getting a clean test of the model longitudinally. In this section we discuss these problems, and discuss some of our results on the longitudinal tests.

A priori we would expect to do a better job making cross-sectional rather than longitudinal predictions. Indeed, it is not clear that this model should predict anything at all about longitudinal variations. To see why, imagine that the assumptions of the model are satisfied perfectly, and suppose that the five parameters of the order flow process (μ, α, etc.) for a given stock are fixed in time. Then the only daily variations we would observe in testing the model would be due to sample errors in the estimation process. Even though the assumptions are satisfied perfectly, we would find no correlation between predicted and actual values. To observe such a correlation requires real variations in the parameters of the order flow process. There are also possible problems with relaxation times: If a parameter is suddenly changed, according to the model it takes the system time to reach a new steady state behavior. There are two characteristic times in the model: σ/μ, which is the characteristic time for removal of limit orders by market orders, and $1/\delta$, which is is the characteristic time for cancellations of limit orders. For the data here it appears that σ/μ is typically less than a minute, whereas $1/\delta$ ranges from a few minutes to a few hours. Thus, $1/\delta$ is the slowest relaxation time, and in some cases at least it is potentially problematic for a daily analysis. In addition, there is the very significant problem that real order flows are strongly autocorrelated, discussed below.

Cross-sectionally, in contrast, we expect *a priori* that different stocks should have different parameters. There are likely to be larger variations in the parameters between stocks than in the parameters for a given stock at different times. In addition, for a cross-sectional analysis there are no problems with relaxation times, and in any case averaging over longer periods of time reduces the sampling error. Thus cross-sectional analysis is expected to be more promising and more reliable.

As noted, for the daily analysis, and even for cross-sectional analysis over long periods of time, there are problems caused by the long range autocorrelations of real order flow, spreads, and price diffusion rates. Autocorrelations can remain strongly positive on the order of 50 days. This creates problems in performing the regression, and can result in a systematic bias in the estimated parameters. It causes severe systematic biases and interpretation problems for a daily analysis.

To produce estimates of the average values of the parameters and of the price diffusion and spread across the full 21 month period for the cross-sectional regressions, we have used the event-weighted average of the daily values. The alternative would have been to repeat the measurements as done for the daily data on a 21 month rather than a daily time-scale. However, this latter approach would run into problems because of the opening auction, which is not treated by our model. There are price changes driven by the orders received during the opening auction, and if we measured price diffusion across the full period we would be including these as well as the intra-day price movements. As a simple solution to this problem we use an event-weighted 21 month average of daily values to compute values for each of the order flow parameters, and then make predictions for each stock based on the average values. The weighting is done by the number of events in a day, which for simple quantities such as the market impact rate reduces to something that is equivalent to applying the analysis over the full period. Similarly, to get the 21 month average of the spread and price diffusion we simply compute an event-weighted average of their daily values. We have tried several variations on this procedure and the differences appear to be inconsequential.

When we perform longitudinal regressions at a daily time-scale we get values for the slope coefficient of the regressions that are less than one, often by a statistically significant amount. We believe this is caused by the strong autocorrelation. For example, consider a time series process of the form

$$y_t = ax_t + \rho y_{t-1} + n_t \tag{3}$$

where n_t is an IID noise process. In case x_t are i.i.d., regressing y_t against x_t will result in coefficients that are systematically too small, due to the fact that the y_{t-1} term damps the response of y_t to changes in x_t. Of course, one can fix this in the simple example above by simply including y_{t-1} in the regression [31]. For the real data, however, the autocorrelation structure is more complicated – indeed we believe it is a long-memory process – which is not well modeled by an AR process in the above form. Without finding a proper characterization of the autocorrelation structure, we are likely to make errors in estimating the dependence of the predicted and actual values. This is borne out in the error analysis presented in Section (4.3), where we see that as we break the data into shorter subsamples, the estimated slope coefficients systematically decrease for the spread and the price diffusion.

5 Average market impact

Market impact is practically important because it is the dominant source of transaction costs for large traders, and conceptually important because it provides a convenient probe of the revealed supply and demand functions as manifested by the limit order book. When a market order of size ω arrives it causes transactions which can cause a change in the midpoint price $m(t) \equiv (a(t) + b(t))/2$. The average market impact function ϕ is the average logarithmic midpoint price shift Δp conditioned on order size, $\phi(\omega) = E[\Delta p|\omega]$.

A long-standing mystery about market impact is that it is highly concave [24, 25, 26, 27, 28, 15, 17, 18]. This is unexpected since simple arguments would suggest that because of the multiplicative nature of returns, market impact should grow at least linearly [5]. The model we are testing predicts a concave average market impact function, with the concavity becoming more pronounced for small ϵ. However, these predictions are not in good detailed agreement with the data; the model predicts a larger variation with ϵ than what is actually observed. This is not surprising, given the assumption of uniform order placement and the subdiffusive nature of prices of the Poisson order flow model mentioned earlier, both of which affect the average market impact.

5.1 Collapse in non-dimensional coordinates

However, a surprising regularity of the average market impact function is uncovered by simply plotting the data in non-dimensional coordinates, as shown in Fig. 9. Each market order ω_i causes a possible price change Δp_i, defining an impact event $(\omega_i, \Delta p_i)$; if the mid price does not change, $\Delta p_i = 0$. If we bin together events with similar ω and plot the mean order size as a function of the mean price impact Δp, we typically see highly variable behavior for different stocks, as shown in Fig. 9(b). We have also explored other ways of renormalizing the order size, such as taking the ratio of each order to the daily or full-sample mean, but they give similar behavior.

Plotting the data in non-dimensional units tells a much simpler story. Here, we normalize each impact event (transaction) by the appropriate normalization constants as measured for the day at which the event occured. The order sizes ω_i we normalize by the characteristic number of shares N_c, and the price impacts δp_i by the characteristic price P_c. Upon binning simmilar ω_i values, the data collapses onto roughly a single curve, as shown in Fig. 9(a). The variations from stock to stock are quite small; on average the corresponding bins for each stock deviate from each other by about 8%, roughly the size of the statistical sampling error. We have made an extensive analysis, but due to problems caused by the long-memory property of these time series and

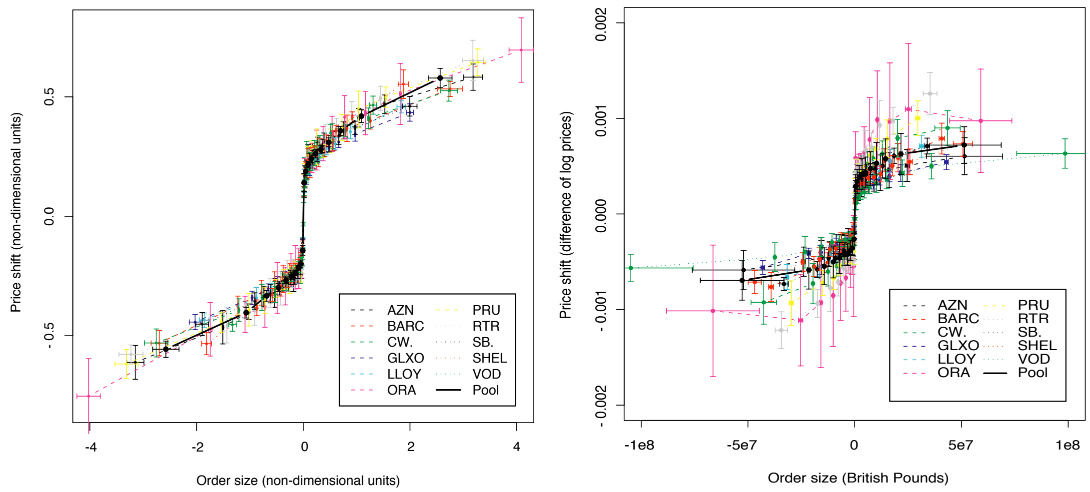

Figure 9: The average market impact as a function of the mean order size. In (a) the price differences and order sizes for each transaction are normalized by the non-dimensional coordinates dictated by the model, computed on a daily basis. Most of the stocks collapse extremely well onto a single curve; there are a few that deviate, but the deviations are sufficiently small that given the long-memory nature of the data and the cross-correlations between stocks, it is difficult to determine whether these deviations are statistically significant (see Appendix 5.3). This means that we understand the behavior of the market impact as it varies from stock to stock by a simple transformation of coordinates. In (b), for comparison we plot the order size in units of British pounds against the average logarithmic price shift.

cross correlations between stocks, it remains unclear whether these differences are statistically significant (see Section 5.3). In contrast, using standard coordinates the differences are highly statistically significant. This collapse illustrates that the non-dimensional coordinates dictated by the model provide substantial explanatory power: We can understand how the average market impact varies from stock to stock by a simple transformation of coordinates. This provides a more fundamental explanation for the empirically constructed collapse of average market impact for the New York Stock Exchange found earlier [17].

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. 10. The functional form of the market impact we observe here is not in agreement with a recent theory by Gabaix et al. [18], 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.

5.2 Market impact vs. supply and demand

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

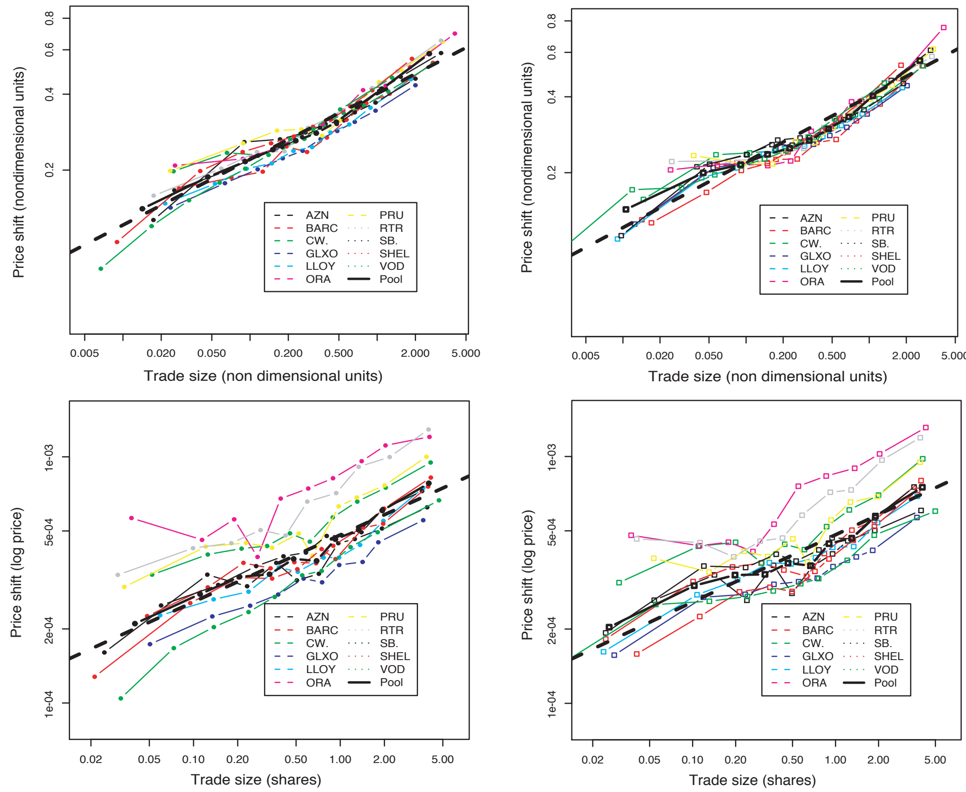

Figure 10: The average market impact vs. order size plotted on log-log scale. The upper left and right panels show buy and sell orders in non-dimensional coordinates; the fitted line has slope $\beta = 0.26 \pm 0.02$ for buy orders and $\beta = 0.23 \pm 0.02$ for sell orders. In contrast, the lower panels 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.

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 [5]. 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.

5.3 Error analysis for market impact

Assigning error bars to the average market impact is difficult because for many stocks, particularly liquid ones, 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 [21, 22]. 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 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 [23], Section 4.4. This is a more straightforward version of the method discussed in Section (4.3). 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 [23] 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 intercepts 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. 11.

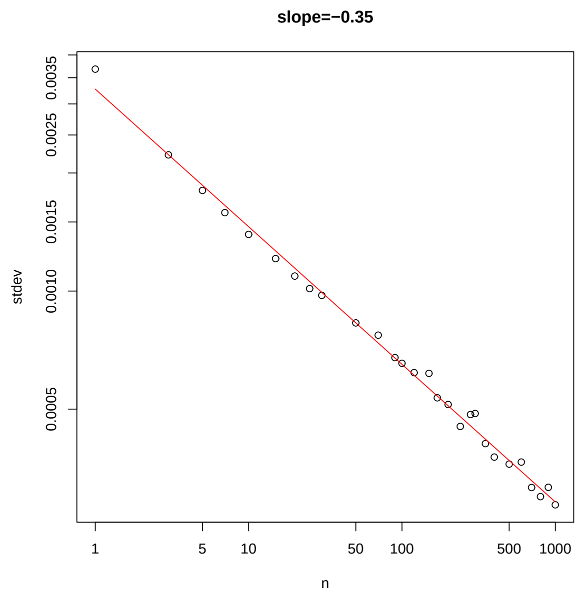

Figure 11: The variance plot procedure used to determine error bars for mean market impact conditional on order size. The horizontal axis n denotes the number of points in the m different samples, and the vertical axis is the standard deviation of the m sample means. We estimate the error of the full sample mean by extrapolating n to the full sample length.

A central question about Fig. 9 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. 9 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 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.

6 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 very well. However, there are several aspects that it does not describe well, such as the scale-free power law properties. This would require a more sophisticated model of order flow, including a more realistic model of price dependence in order placement and cancellations [15, 20], long-memory properties [21, 22] 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. While this would have some advantages, it would also have some disadvantages.

Some market properties that might profit from such an improved model are detailed below.

- *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, with

H close to and typically a little greater than 0.5. In contrast, under Poisson order flow, due to the dynamics of the double continuous auction price formation process, prices make a strongly anti-correlated random walk, so 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. [4, 5] showed that the limits $\tau \to 0$ and $\tau \to \infty$ 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 $D(0)$ does a good job of matching the real data, as we have shown here, while $D(\infty)$ does a poor job. Note that it is very interesting that the double continuous auction produces anti-correlations in prices, even with no correlation in order flow. One can turn this around: Given that prices are uncorrelated, there must be correlations in order flow. And indeed this is observed to be the case [21, 22].

- *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 (4.4) the problems that the autocorrelations in spread and price diffusion create for comparing the theory to the model on a daily scale.
- *Lack of dependence on granularity parameter.* In Section (5.3) 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.

7 Conclusions

The model we have presented here does a strikingly good job of predicting the average spread, and a pretty good job of predicting the price diffusion rate. Also, by simply plotting the data in non-dimensional coordinates we get a good prediction of market impact. These are remarkable because to do this we have completely dropped agent rationality, instead focusing all our attention on the problem of understanding the constraints imposed by the continuous double auction. The resulting dynamics of the interaction between price formation and order placement are complicated, and the predictions of the model are nontrivial.

The approach taken here can be viewed as a divide and conquer strategy. Rather than attempting to explain the properties of the market from fundamental assumptions about utility maximization by individual agents, we divide the problem into two parts. The first problem, which is addressed here, is that of understanding the characteristics of the market given a model of order flow. The second (and harder) problem, which remains to be investigated, is that of explaining order flow. This is part of a broader research program that might be characterized as the "low-intelligence" approach to economics: We begin with zero-intelligence agents to get a good benchmark of the effect of market institutions, and once this benchmark is well-understood, add a little intelligence, moving toward market efficiency. We thus start from zero rationality and work our way up, in contrast to the canonical approach of starting from perfect rationality and working down. See Ref. [30].

These results have several practical implications. For market practitioners, understanding the spread and the market impact function is very useful for estimating transaction costs and for developing algorithms that minimize their effect. And for regulators they suggest that it may be possible to make prices less volatile and lower transaction costs by creating incentives for limit orders and disincentives for market orders.

The model we test here was constructed before looking at the data [4, 5], and was designed to be as simple as possible for analytic analysis. A more realistic (but necessarily more complicated) model would more closely mimic the properties of real order flows, which are price dependent and strongly correlated both in time and across price levels. An improved model would hopefully be able to capture even more features of the data, such as the power law tails of prices. Nonetheless, as we have shown above, this simple model does a remarkable job of explaining important fundamental properties of markets, such as transaction costs, price diffusion and supply and demand. The model captures the statistical mechanics of the market quite well, and in particular, the way order placement and price formation interact to alter the accumulation of stored supply and demand. For the phenomena studied here this appears to be the dominant effect. We do not mean to claim that market participants are unintelligent: 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 strategic behavior. It is surprising that such a simple model can explain so much about a system as complex as a market, and shed light on century-old questions about the rate of price diffusion and the form of supply and demand

References and acknowledgments

[1] L. Bachelier, "Théorie de la spéculation", 1900. Reprinted in P.H. Cootner, *The Random Character of Stock Prices*, 1964, MIT Press Cambridge.

[2] Becker, G., Irrational behavior and economic theory, *J. of Political Economy* **70**, 1 (1962) 1-13.

[3] D. Gode and S. Sunder, Allocative efficiency of markets with zero intelligence traders: Markets as a partial substitute for individual rationality. *J. of Political Economy*, **101**, 119 (1993).

[4] Daniels, M, Farmer, J.D., Gillemot, L., Iori, G., and Smith, D.E. Quantitative model of price diffusion and market friction based on trading as a mechanistic random process, *Physical Review Letters* **9019**, 10 2003.

[5] Smith, E., Farmer, J.D., Gillemot, L., and Krishnamurthy, S., Statistical theory of the continuous double auction, To appear in *Quantitative Finance*, 2003.

[6] H. Mendelson, Market behavior in a clearing house, *Econometrica* **50**, 1505-1524 (1982).

[7] 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. Schwartz, eds., *Market Making and the Changing Structure of the Securities Industry*, 1985, Lexington Books, Lexington MA.

[8] I. Domowitz and Jianxin Wang, Auctions as algorithms, *J. of Econ. Dynamics and Control* **18**, 29 (1994).

[9] T. Bollerslev, I. Domowitz, and J. Wang, Order flow and the bid-ask spread: An empirical probability model of screen-based trading, *J. of Econ. Dynamics and Control* **21**, 1471 (1997).

[10] P. Bak, M. Paczuski, and M. Shubik, Price variations in a stock market with many agents, *Physica A* **246**, 430 (1997).

[11] D. Eliezer and I.I. Kogan, Scaling laws for the market microstructure of the interdealer broker markets, http://xxx.lanl.gov/cond-mat/9808240.

[12] S. Maslov, Simple model of a limit order-driven market, *Physica A* **278**, 571(2000).

[13] F. Slanina, Mean-field approximation for a limit order driven market model, *Phys. Rev. E*, **64**, 056136 (2001).

[14] D. Challet and R. Stinchcombe, Analyzing and Modeling 1+1d markets, *Physica A* **300**, 285, (2001).

[15] J.-P. Bouchaud, M. Mezard, M. Potters, Statistical properties of the stock order books: empirical results and models, *Quantitative Finance* **2** (2002) 251-256.

[16] Potters, M. & Bouchaud, J.-P., "More statistical properties of orders books and price impact", *Physica A* **324**, 133-140 (2003)

[17] Lillo,F. Farmer. J.D. & Mantegna, R.N., Master Curve for Price-Impact Function, *Nature* **421**, 129-130 (2003).

[18] Gabaix, X., Gopikrishnan, P., Plerou, V. and Stanley, H.E. A theory of power-law distributions in financial market fluctuations, *Nature* **423**, 267-270 (2003).

[19] *SETS four years on - October 2001*, published by the London Stock Exchange

[20] Zovko, I. and Farmer, J.D., The power of patience: A behavioral regularity in limit order placement, *Quantitative Finance*, October 2002.

[21] Bouchaud,J.-P., Gefen,Y., Potters,M., and Wyart,M. (2003). Fluctuations and response in financial markets: the subtle nature of 'random' price changes, xxx.lanl.gov/cond-mat/0307332.

[22] Lillo, F. and Farmer, J.D., The long memory of the efficient market, xxx.lanl.gov/cond-mat/0311053 (2003).

[23] Beran, J. *Statistics for Long-Memory Processes*, Chapman & Hall (1994).

[24] J.A. Hausman and A.W. Lo, "An ordered probit analysis of transaction stock prices", *Journal of Financial Economics* **31** 319-379 (1992)

[25] J.D. Farmer, "Slippage 1996", Prediction Company internal technical report (1996), http://www.predict.com/jdf/slippage.pdf

[26] N. Torre, *BARRA Market Impact Model Handbook*, BARRA Inc, Berkeley CA, www.barra.com (1997).

[27] A. Kempf and O. Korn, "Market depth and order size", University of Mannheim technical report (1998).

[28] V. Plerou, P. Gopikrishnan, X. Gabaix, and H.E. Stanley, "Quantifying stock price response to demand fluctuations", Phys. Rev. E **66** 027104 (2002).

[29] Farmer, J.D. and Lillo, F., On the origin of power law tails in price fluctuations, xxx.lanl.gov/cond-mat/0309416.

[30] Nelkin, I., Innovations in trading strategies, *Quantitative Finance* **3**, Number 4, C63-74 (2003). (See relevant remarks by J.D. Farmer).

[31] William H. Greene, Econometric Analysis, Prentice Hall.

We would like to thank Credit Suisse First Boston, the James S. McDonnell Foundation, the McKinsey Corporation, Bob Maxfield, and Bill Miller for their support of this project. We would also like to thank Sam Bowles, Supriya Krishnamurthy, Fabrizio Lillo, and Eric Smith for useful discussions, and Mark Bieda, David Krakauer, Harold Morowitz and Elizabeth Wood for comments on the manuscript.