QUANTITATIVE FINANCE VOLUME 2 (2002) 387–392 RESEARCH PAPER
INSTITUTE OF PHYSICS PUBLISHING quant.iop.org

The power of patience: a behavioural regularity in limit-order placement

Ilija Zovko[1,2] and J Doyne Farmer[1]

[1] Santa Fe Institute, 1399 Hyde Park Rd, Santa Fe, NM 87501, USA
[2] CeNDEF, University of Amsterdam, Roetersstraat 11, Amsterdam, The Netherlands

E-mail: zovko@santafe.edu and jdf@santafe.edu

Received 19 June 2002, in final form 11 September 2002
Published 27 September 2002
Online at stacks.iop.org/Quant/2/387

Abstract

In this paper we demonstrate a striking regularity in the way people place limit orders in financial markets, using a data set consisting of roughly two million orders from the London Stock Exchange. We define the relative limit price as the difference between the limit price and the best price available. Merging the data from 50 stocks, we demonstrate that for both buy and sell orders, the unconditional cumulative distribution of relative limit prices decays roughly as a power law with exponent approximately −1.5. This behaviour spans more than two decades, ranging from a few ticks to about 2000 ticks. Time series of relative limit prices show interesting temporal structure, characterized by an autocorrelation function that asymptotically decays as $C(\tau) \sim \tau^{-0.4}$. Furthermore, relative limit price levels are positively correlated with and are led by price volatility. This feedback may potentially contribute to clustered volatility.

1. Introduction

Most modern financial markets are designed as a complex hybrid composed of a continuous double auction and an ‘upstairs’ trading mechanism serving the purpose of block trades. The double auction is believed to be the primary price discovery mechanism[3]. Limit orders, which specify both a quantity and a limit price (the worst acceptable price), are the liquidity-providing mechanism for double auctions and the proper understanding of their submission process is important in the study of price formation.

We study the *relative limit price* $\delta(t)$, the limit price in relation to the current best price. For buy orders $\delta(t) = b(t) - p(t)$, where p is the limit price, b is the best bid (highest buy limit price) and t is the time when the order is placed. For sell orders $\delta(t) = p(t) - a(t)$, where a is the best ask (lowest sell limit price)[4]. We find a striking regularity in the distribution of relative limit prices and we document clustering of order prices as seen by a slowly decaying autocorrelation function.

Biais *et al* (1995) studied the limit-order submission process on the Paris Bourse. They note that the number of orders placed up to five quotes away from the market decays monotonically but do not attempt to estimate the distribution or examine orders placed further than five best quotes. Our analysis looks at the price placement of limit orders across a much wider range of prices. Since placing orders out of the

[3] According to the London Stock Exchange information bulletins (‘SETS four years on—October 2001’, published by the London Stock Exchange), since the introduction of the SETS in 1997 to October 2001, the average percentage of trades in order book securities that have been executed at the price shown on the order book is 70–75%. Therefore SETS seems to serve as the primary price discovery mechanism in London.

[4] We have made a somewhat arbitrary choice in defining the reference price. An obvious alternative would have been to choose the best ask as the reference price for buy orders and the best bid as the reference price for sell orders. This would have the advantage that it would have automatically included orders placed inside the interval between the bid and ask (the spread), which are discarded in the present analysis. However, the choice of reference price does not seem to make a large difference in the tail; for large δ it leads to results that are essentially the same.

market carries execution and adverse selection risk, our work is relevant in understanding the fundamental dilemma of limit-order placement: execution certainty versus transaction costs (see, e.g., Cohen *et al* 1981, Harris 1997, Harris and Hasbrouck 1996, Holden and Chakravarty 1995, Kumar and Seppi 1992, Lo *et al* 2002).

In addition to the above, our work relates to the literature on clustered volatility. It is well known that both asset prices and quotes display ARCH/GARCH effects (Engle 1982, Bollerslev 1986), but the origins of these phenomena are not well understood. Explanations range from news clustering (Engle *et al* 1990), macroeconomic origins (Campbell 1987, Glosten *et al* 1993) to microstructure effects (Lamoureux and Lastrapes 1990, Bollerslev and Domowitz 1991, Kavajecz and Odders-White 2001). We provide empirical evidence that volatility feedback may in part be caused by limit-order placement that in turn depends on past volatility levels.

This paper is organized as follows. Section 2 introduces the mechanics of limit-order trading and describes the London Stock Exchange data we use. Section 3 presents our results on the distribution and time series properties of relative limit-order prices. In section 4 we examine the possible relationship of limit-order prices and volatility which may lead to volatility clustering. Section 5 discusses and summarizes the result.

2. Description of the London Stock Exchange data

The limit-order trading mechanism works as follows: as each new limit order arrives, it is matched against the queue of pre-existing limit orders, called the *limit-order book*, to determine whether or not it results in any immediate transactions. At any given time there is a best buy price $b(t)$ and a best ask price $a(t)$. A sell order that crosses $b(t)$, or a buy order that crosses $a(t)$, results in at least one transaction. The matching for transactions is performed based on price and order of arrival. Thus matching begins with the order of the opposite sign that has the best price and arrived first, then proceeds to the order (if any) with the same price that arrived second, and so on, repeating for the next best price, etc. The matching process continues until the arriving order has either been entirely transacted, or until there are no orders of the opposite sign with prices that satisfy the arriving order's limit price. Anything that is left over is stored in the limit-order book.

A crossing limit order is defined as a limit order that results in at least a partial immediate transaction. Traders submit such orders to limit their market impact. Crossing limit orders make up about 30% (in the example of Vodafone) of all limit orders and are more like market orders. In this paper we discard them and analyse only limit orders that enter the book. Of the analysed orders 74% are submitted at the best quotes. Only 1% are submitted inside the spread (with $\delta < 0$), while the remaining 25% are submitted out of the market ($\delta > 0$). We investigate only limit orders with positive relative price $\delta > 0$ and refer to them in the text simply as limit orders[5].

The time period of the analysis is from 1 August 1998 to 31 April 2000. This data set contains many errors; we chose the names we analyse here from the several hundred that are traded on the exchange based on the ease of cleaning the data, trying to keep a reasonable balance between high and low volume stocks[6]. This left 50 different names, with a total of roughly seven million limit orders, of which about two million are submitted out of the market ($\delta > 0$).

3. Properties of relative limit-order prices

Choosing a relative limit price is a strategic decision that involves a trade-off between patience and profit (see, e.g., Holden and Chakravarty 1995, Harris and Hasbrouck 1996, Sirri and Peterson 2002). Consider a sell order; the story for buy orders is the same, interchanging 'high' and 'low'. An impatient seller will submit a limit order with a limit price well below $b(t)$, which will immediately result in a transaction. A seller of intermediate patience will submit an order with $p(t)$ a little greater than $b(t)$; this will not result in an immediate transaction, but will have high priority as new buy orders arrive. A very patient seller will submit an order with $p(t)$ much greater than $b(t)$. This order is unlikely to be executed soon, but it will trade at a good price if it does. A higher price is clearly desirable, but it comes at the cost of lowering the probability of trading—the higher the price, the lower the probability there will be a trade. The choice of limit price is a complex decision that depends on the goals of each agent. There are many factors that could affect the choice of limit price, such as the time horizon of the trading strategy. *A priori* it is not obvious that the unconditional distribution of limit prices should have any particular simple functional form.

3.1. Unconditional distribution

Figure 1 shows examples of the cumulative distribution for stocks with the largest and smallest numbers of limit orders. Each order is given the same weighting, regardless of the number of shares, and the distribution for each stock is normalized so that it sums to one. There is considerable variation in the sample distribution from stock to stock, but these plots nonetheless suggest that power-law behaviour for large δ is a reasonable hypothesis. This is somewhat clearer for the stocks with high order arrival rates. The low volume stocks show larger fluctuations, presumably because of their smaller sample sizes. Although there is a large number of events in each of these distributions, as we will show later, the samples are highly correlated, so that the effective number of independent samples is not nearly as large as it seems.

[5] Even though limit orders placed in the spread are not numerous, they are very important in price formation. The data set we use does not include enough events to provide statistically significant results for such orders. Our preliminary results indicate that orders placed in the spread behave qualitatively similar to orders placed out of the market, i.e. there are some indications of power-law behaviour in their limit price density towards the other side of the market.

[6] The ticker symbols for the stocks in our sample are AIR, AL., ANL, AZN, BAA, BARC, BAY, BLT, BOC, BOOT, BPB, BSCT, BSY, BT.A, CCH, CCM, CS., CW., GLXO, HAS, HG., ICI, III, ISYS, LAND, LLOY, LMI, MKS, MNI, NPR, NU., PO., PRU, PSON, RB., RBOS, REED, RIO, RR., RTK, RTO, SB., SBRY, SHEL, SLP, TSCO, UNWS, UU., VOD, and WWH.

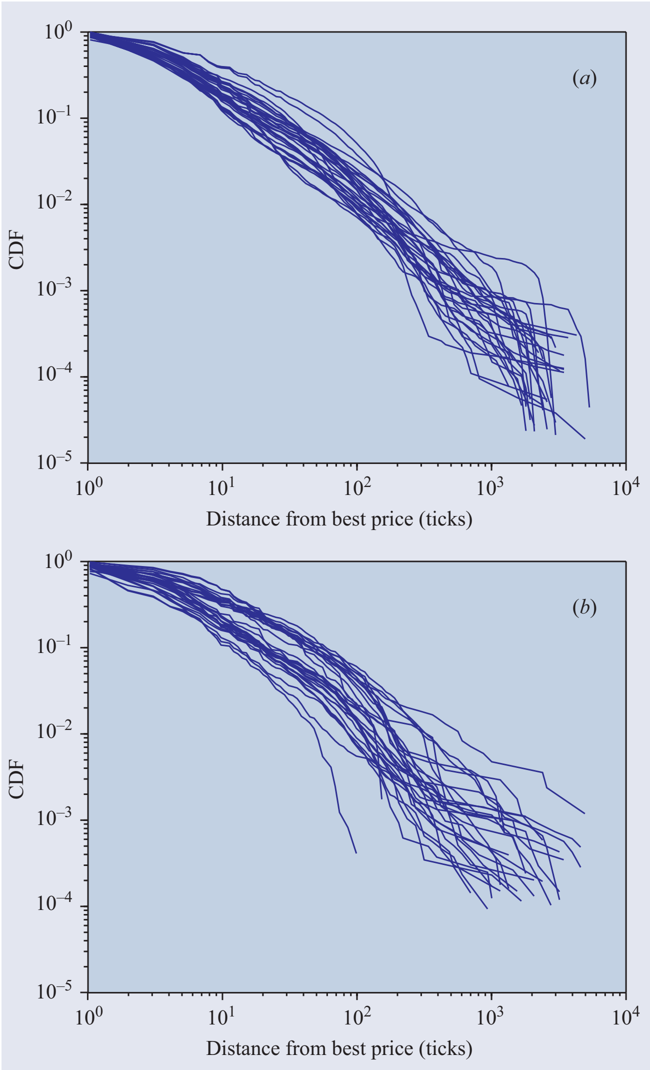

Figure 1. (a) Cumulative distribution functions $P(\delta) = \text{prob}\{x \geqslant \delta\}$ of relative limit price δ for both buy and sell orders for the 15 stocks with the largest number of limit orders during the period of the sample (those that have between 150 000 and 400 000 orders). (b) The same for 15 stocks with the lowest number of limit orders, in the range 2000–100 000. (To avoid overcrowding, we have averaged together nearby bins, which is why the plots appear to violate the normalization condition.)

To reduce the sampling errors we merge the data for all stocks, and estimate the sample distribution for the merged set using the method of ranks, as shown in figure 2. We fit the resulting distribution to the functional form[7]

$$P(\delta) = \frac{A}{(x_0 + \delta)^\beta}. \tag{1}$$

[7] The functional form we use to fit the distribution has to satisfy two requirements: it has to be a power law for large δ and finite for $\delta = 0$. A pure power law is either not integrable at 0 or at ∞. If the functional form is to be interpreted as a probability density then it necessarily has to be truncated at one end. In our case the natural truncation point is 0. Clearly there is some arbitrariness in the choice of the exact form, but since we are mainly interested in the behaviour for large δ, this functional form seems satisfactory.

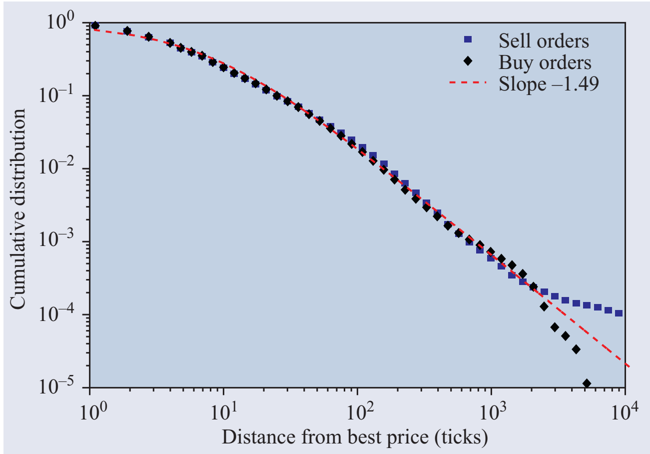

Figure 2. An estimate of the cumulative probability distribution based on a merged data set, containing the relative limit-order sizes $\delta(t)$ for all 50 stocks across the entire sample. The solid curve is a nonlinear least squares fit to the logarithmic form of equation (1).

A is set by the normalization and is a simple function of x_0 and β. Fitting this to the entire sample (both buys and sells) gives $x_0 = 7.01 \pm 0.05$ and $\beta = 1.491 \pm 0.001$. When treated separately, buys and sells gave similar values for the exponent, i.e. $\beta = 1.49$ in both cases. Since these error bars based on goodness of fit are certainly overly optimistic, we also tested the stability of the results by fitting buys and sells separately on the first and last halves of the sample, which gave values in the range $1.47 < \beta < 1.52$. Furthermore, we checked whether there are significant differences in the estimated parameters for stocks with high versus low order arrival rates. The results ranged from $\beta = 1.5$ for high to $\beta = 1.7$ for low arrival rates, but for the low arrival rate group we do not have high confidence in the estimate.

As one can see from the figure, the fit is reasonably good. The power law is a good approximation across more than two decades, for relative limit prices ranging from about 10–2000 ticks. For British stocks, ticks are measured either in pence, half pence or quarter pence; in the former case, 2000 ticks corresponds to about 20 pounds. Given the low probability of execution for orders with such high relative limit prices this is quite surprising. (For Vodafone, for example, the highest relative limit price that eventually resulted in a transaction was 240 ticks.) The value of the exponent $\beta \approx 1.5$ implies that the mean of the distribution exists, but its variance is formally infinite. Note that because normalized power-law distributions are scale free, the asymptotic behaviour does not depend on units, e.g. ticks versus pounds. There appears to be a break in the power law at about 2000 ticks, with sell orders deviating above and buy orders deviating below. A break at roughly this point is expected for buy orders due to the fact that $p = 0$ places a lower bound on the limit price. For a stock trading at 10 pounds, for example, with a tick size of half a pence, 2000 ticks is the lowest possible relative limit price for a buy order. The reason for a corresponding break for sell orders is not so obvious, but in view of the extremely low probability of execution, is not surprising. It should also be kept in mind that the number of events in the extreme tail is very low, so this could also be a statistical fluctuation.

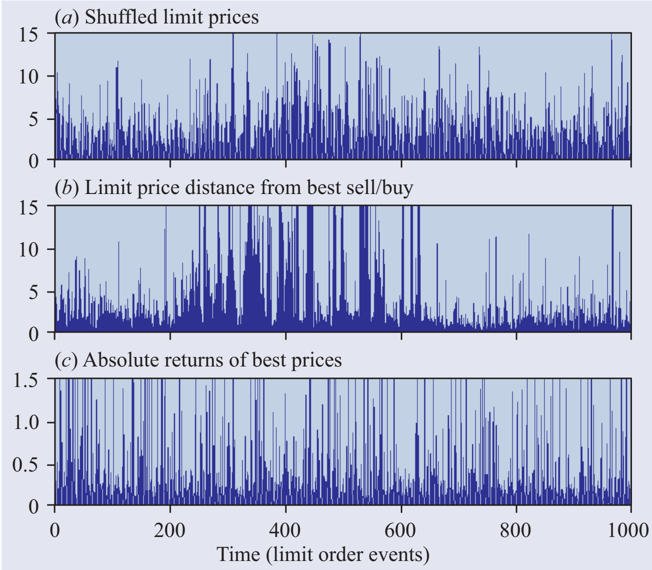

Figure 3. (a) Time series of randomly shuffled values of $\delta(t)$ for stock Barclays Bank. (b) True time series $\delta(t)$. (c) The absolute value of the change in the best price between each event in the $\delta(t)$ series.

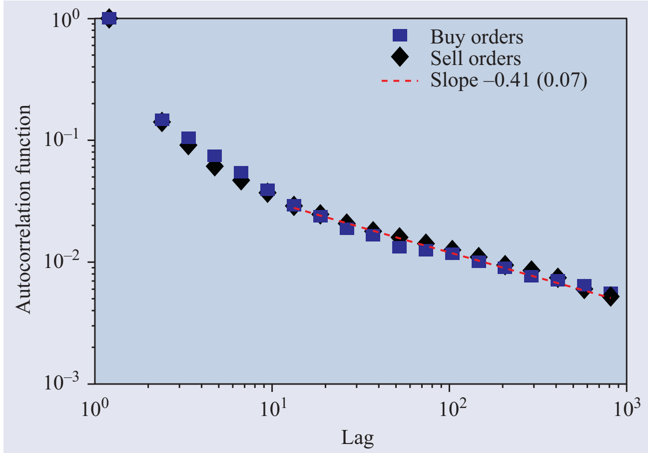

Figure 4. The autocorrelation of the time series of relative limit prices δ, averaged across all 50 stocks in the sample and smoothed across different lags. This is computed in tick time, i.e. the x axis indicates the number of events, rather than a fixed time.

3.2. Time series properties

The time series of relative limit prices also has interesting temporal structure. This is apparent to the eye, as seen in figure 3(b), which shows the average relative limit price $\bar{\delta}$ in intervals of approximately 60 events for Barclays Bank. For reference, in figure 3(a) we show the same series with the order of the events randomized. Comparing the two suggests that the large and small events are more clustered in the real series than in the shuffled series.

This temporal structure appears to be described by a slowly decaying autocorrelation function, as shown in figure 4. Since the second moment of the unconditional distribution does not exist, there are potential problems in computing the autocorrelation function. The standard deviations in the denominator formally do not exist, and the terms in the numerator can be slow to converge. To cope with this we have imposed a cut-off at 1000 ticks, averaged across all 50 stocks in our sample, and smoothed the autocorrelation function for large lags (where the statistical significance drops). The resulting average autocorrelation function decays asymptotically as a power law of the form $C(\tau) \sim \tau^{-\gamma}$, with $\gamma \approx 0.4$, indicating that relative limit price placement is quite persistent with no characteristic timescale. In the figure, we have computed the autocorrelation function in tick time, i.e. the lags correspond to the event order. This means that low order arrival volume stocks have longer real time intervals than high order arrival volume stocks. We have also obtained a similar result using real time, by computing the mean limit price $\bar{\delta}$ in 13 min intervals (merging daily boundaries). In this case the behaviour is not quite as regular but is still qualitatively similar. We still see a slowly decaying power-law tail, though with a somewhat lower exponent (roughly 0.3). The autocorrelations are quite significant even for lags of 1000, corresponding to about eight days. Roughly the same behaviour is seen for buy and sell orders, and for the first ten months and the last ten months of the sample. We computed error bars for this result by randomly shuffling the time series 100 times, and computing the 2.5 and 97.5% quantiles of the sample autocorrelation for each lag. This gives error bars at roughly $\pm 10^{-3}$.

One consequence of such a slowly decaying autocorrelation is the slow convergence of sample distributions to their limiting distribution. If we generate artificial IID data with equation (1) as the unconditional distribution, the sample distributions converge very quickly with only a few thousand points. In contrast, for the real data, even for a stock with 200 000 points, the sample distributions display large fluctuations. When we examine subsamples of the real data, the correlations in the deviations across subsamples are obvious and persist for long periods of time, even when there is no overlap in the subsamples. We believe that the slow convergence of the sample distributions is mainly due to the long-range temporal dependence in the data.

4. Volatility clustering

To get some insight into the possible cause of the temporal correlations, we compare the time series of relative limit prices to the corresponding price volatility. The price volatility is measured as $v(t) = |\log(b(t)/b(t-1))|$, where $b(t)$ is the best bid for buy orders or the best ask for sell limit orders. We show a typical volatility series in figure 3(c). One can see by eye that epochs of high limit price tend to coincide with epochs of high volatility.

To help understand the possible relation between volatility and relative limit price we calculate their cross-autocorrelation. This is defined as

$$XCF(\tau) = \frac{\langle v(t-\tau)\delta(t)\rangle - \langle v(t)\rangle\langle\delta(t)\rangle}{\sigma_v \sigma_\delta}, \qquad (2)$$

where $\langle\cdot\rangle$ denotes a sample average and σ denotes the standard deviation. We first create a series of the average relative limit price and average volatility over 10 min intervals. We then

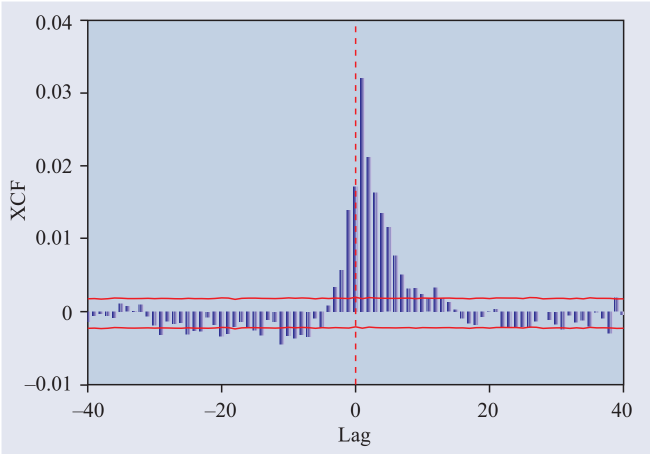

Figure 5. The cross autocorrelation of the time series of relative limit prices $\delta(t)$ and volatilities $v(t-\tau)$, averaged across all 50 stocks in the sample.

compute the cross-autocorrelation function and average over all stocks. The result is shown in figure 5.

We test the statistical significance of this result by testing against the null hypothesis that the volatility and relative limit price are uncorrelated. To do this we have to cope with the problem that the individual series are highly autocorrelated, as demonstrated in figure 4, and the 50 series for each stock also tend to be correlated to each other. To solve these problems, we construct samples of the null hypothesis using a technique introduced in Theiler *et al* (1992). We compute the discrete Fourier transform of the relative limit price time series. We then randomly permute the phases of the series and perform the inverse Fourier transform. This creates a realization of the null hypothesis, drawn from a distribution with the same unconditional distribution and the same autocorrelation function. Because we use the same random permutation of phases for each of the 50 series, we also preserve their correlation with each other. We then compute the cross-autocorrelation function between each of the 50 surrogate limit price series and its corresponding true volatility series, and then average the results. We then repeat this experiment 300 times, which gives us a distribution of realizations of averaged sample cross-autocorrelation functions under the null hypothesis. This procedure is more appropriate in this case than the standard moving block bootstrap, which requires choice of a timescale and will not work for a series such as this that does not have a characteristic timescale. The 2.5 and 97.5% quantile error bars at each lag are denoted by the two solid lines near zero in figure 5.

From this figure it is clear that there is indeed a strong contemporaneous correlation between volatility and relative limit price, and that the result is highly significant. Furthermore, there is some asymmetry in the cross-autocorrelation function; the peak occurs at a lag of one rather than zero and there is more mass on the right than on the left. This suggests that there is some tendency for volatility to lead the relative limit price. This implies one of three things:

(1) volatility and limit price have a common cause, but this cause is for some reason felt later for the relative limit price;
(2) the agents placing orders key off volatility and correctly anticipate it; or, more plausibly,
(3) volatility at least partially causes the relative limit price. Angel (1994) has suggested that volatility might affect limit-order placement in this way.

Note that this suggests an interesting feedback loop: holding other aspects of the order placement process constant, an increase in the average relative limit price will lower the depth in the limit-order book at any particular price level and therefore increase volatility. Since such a feedback loop is unstable, there are presumably nonlinear feedbacks of the opposite sign that eventually damp it. Nonetheless, such a feedback loop may potentially contribute to creating clustered volatility.

5. Conclusion

One of the most surprising aspects of the power-law behaviour of relative limit price is that traders place their orders so far away from the current price. As is evident in figure 2, orders occur with relative limit prices as large as 10 000 ticks (or 25 pounds for a stock with ticks in quarter pence). While we have taken some precautions to screen for errors, such as plotting the data and looking for unreasonable events, despite our best efforts, it is likely that there are still data errors remaining in this series. There appears to be a break in the merged unconditional distribution at about 2000 ticks; if this is statistically significant, it suggests that the very largest events may follow a different distribution from the rest of the sample, and might be dominated by data errors. Nonetheless, since we know that most of the smaller events are real, and since we see no break in the behaviour until roughly $\delta \approx 2000$, errors are highly unlikely to be the cause of the power-law behaviour seen for $\delta < 2000$.

The conundrum of very large limit orders is compounded by consideration of the average waiting time for execution as a function of relative limit price. We intend to investigate the dependence of the waiting time on the limit price in the future, but since this requires tracking each limit order, the data analysis is more difficult. We have checked this for one stock, Vodafone, in which the largest relative limit price that resulted in an eventual trade was $\delta = 240$ ticks. Assuming other stocks behave similarly, this suggests that either traders are strongly over-optimistic about the probability of execution or that the orders with large relative limit prices are placed for other reasons.

Since obtaining our results, we have seen a recent preprint by Bouchaud *et al* (2002) analysing three stocks on the Paris Bourse over a period of a month. They also obtain a power law for $P(\delta)$, but they observe an exponent $\beta \approx 0.6$, in contrast to our value $\beta \approx 1.5$. We do not understand why there should be such a discrepancy in results. While they analyse only three stock-months of data, whereas we have analysed roughly 1050 stock-months, their order arrival rates are roughly 20 times higher than ours, and their sample distributions appear to follow the power-law scaling fairly well.

One possible explanation is the long-range correlation. Assuming the Paris data show the same behaviour we have observed, the decay in the autocorrelation is so slow that there may not be good convergence in a month, even with a large number of samples. The sample exponent $\hat{\beta}$ based on one month samples may vary with time, even if the sample distributions appear to be well-converged. It is of course also possible that the French behave differently from the British, and that for some reason the French prefer to place orders much further from the midpoint.

Our original motivation for this work was to model price formation in the limit-order book, as part of the research programme for understanding the volatility and liquidity of markets outlined in Daniels *et al* (2001). $P(\delta)$ is important for price formation, since where limit orders are placed affects the depth of the limit-order book and hence the diffusion rate of prices. The power-law behaviour observed here has important consequences for volatility and liquidity that will be described in a future paper.

Our results here are interesting for their own sake in terms of human psychology. They show how a striking regularity can emerge when human beings are confronted with a complicated decision problem. Why should the distribution of relative limit prices be a power law, and why should it decay with this particular exponent? Our results suggest that the volatility leads the relative limit price, indicating that traders probably use volatility as a signal when placing orders. This supports the obvious hypothesis that traders are reasonably aware of the volatility distribution when placing orders, an effect that may contribute to the phenomenon of clustered volatility. Plerou *et al* (1999) have observed a power law for the unconditional distribution of price fluctuations. It seems that the power law for price fluctuations should be related to that of relative limit prices, but the precise nature and the cause of this relationship is not clear. The exponent for price fluctuations of individual companies reported by Plerou *et al* is roughly 3, but the exponent we have measured here is roughly 1.5. Why these particular exponents? Makoto Nirei has suggested that if traders have power-law utility functions, under the assumption that they optimize this utility, it is possible to derive an expression for β in terms of the exponent of price fluctuations and the coefficient of risk aversion. However, this explanation is not fully satisfying, and more work is needed. At this point the underlying cause of the power-law behaviour of relative limit prices remains a mystery.

Acknowledgments

We would like to thank Makoto Nirei, Paolo Patelli, Eric Smith and Spyros Skouras for valuable conversations and Marcus Daniels for valuable technical support. We also thank the McKinsey Corporation, Credit Suisse First Boston, Bob Maxfield and Bill Miller for supporting this research.

References

Angel J 1994 Limit vs. market orders, unpublished working paper

Biais B, Hillion P and Spatt C 1995 An empirical analysis of the limit order book and the order flow in the Paris Bourse *J. Finance* **50** 1655–89

Bollerslev T 1986 Generalized autoregressive conditional heteroskedasticity *J. Econometrics* **31** 307–27

Bollerslev T and Domowitz I 1991 Price volatility, spread variability, and the role of alternative market mechanisms *Rev. Futures Markets* **10** 79–102

Bouchaud J-P, Mezard M and Potters M 2002 Statistical properties of the stock order books: empirical results and models *Quantitative Finance* **2** 251–6

Daniels M, Farmer J D, Iori G and Smith D E 2001 How storing supply and demand affects price diffusion *Preprint* http://xxx.lanl.gov/cond-mat/0112422

Engle R 1982 Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation *Econometrica* **50** 987–1007

Harris L and Hasbrouck J 1996 Market vs. limit orders: the SuperDOT evidence on order submission strategy *J. Financial Quant. Anal.* **31** 213–31

Holden C and Chakravarty S 1995 An integrated model of market and limit orders *J. Financial Intermediation* **4** 213–41

Kavajecz K and Odders-White E 2001 Volatility and market structure *J. Financial Markets* **4** 359–84

Lo A, MacKinlay C and Zhang J 2002 Econometric models of limit-order executions *J. Financial Economics* **65** 31–71

Nirei M, private communication.

Peterson M and Sirri E 2002 Order submission strategy and the curious case of marketable limit orders *J. Financial Quant. Anal.* **37** 221–41

Plerou V, Gopikrishnan P, Amaral L A N, Gabaix X and Stanley H E 1999 Scaling of the distribution of price fluctuations of individual companies *Preprint* http://xxx.lanl.gov/cond-mat/9907161.

Theiler J, Galdrikian B, Longtin A, Eubank S and Farmer J D 1992 Detecting nonlinear structure in time series *Physica* D **58** 77–94