Statistical Regularities in Limit Order Placement

Michael Twardos

Center for Nonlinear Studies, Los Alamos National Lab, New Mexico
87545, USA & The Santa Fe Institute, Santa Fe, New Mexico 87507

J. Doyne Farmer

The Santa Fe Institute, Santa Fe, New Mexico 87507
(Dated: March 14, 2007)

We provide an extensive review of the statistical properties of limit order placements (orders which do not immediately transact) for four stocks on the London Stock Exchange. We look at limit order placements in the context of the distance from the current same best and explore the probability distributions in units of both ticks and logarithms (percentages). We review how orders placed within the spread need to be weighted by the probability that such a limit order could be placed and explore how this changes the probability distribution. We also look at a symmetry measure of weighted and unweighted distributions in both ticks and logarithms to explore the hypothesis that the placements are symmetric around same best. We conclude by reviewing the affect that other quantities such as volatility, price change and spread have on the placement of limit orders.

I. BACKGROUND

Price formation in a double auction market is determined by the placement and cancellation of orders. Ultimately, the decisions of traders to place and cancel orders determines prices changes, volatility and liquidity within the market. In double auction markets, a trader is given the choice of placing a market order that transacts at opposite best or a limit order for any value below opposite best. Between the opposite best and the same best is the spread. The order can be placed within the spread for relatively aggressive traders or outside the spread for more patient traders.

There have been some work on limit order placement for orders less than the same best on the Paris Stock Exchange [1] and the London Stock Exchange [2]. These distributions had a power law form with $\alpha \approx 0.8$ in the PSE and $\alpha \approx 1.5$ in the LSE. The small values of the exponent are surprising because they imply a nonvanishing probability for order placement even at prices that are extremely far from the current best prices, where it would seem that the probability of making a transaction is extremely low. Other studies have tried to understand these problems in terms of a utility maximization problem [3]

Understanding the statistics of limit order placement is a useful way to quantify behavioral regularities in the market. Although, the actions of investors appear to succumb to the whim of personal perceptions of an asset's value, reaction to news events or stock prices, a functional form of placement is a clear contradiction to this perception. Using this information, one may be able to construct a "zero intelligence" model of the market [4]. By creating a model in which orders are randomly chosen from a distribution, we may be able to closely describe price fluctuations and market characteristics observed in the real world. By only using a few parameters, such as the exponent for limit order placement, we may be able to describe market behavior with a closed form equation of state.

II. DETAIL OF THE MEASUREMENT

In this paper, we focus only on orders occurring during the regular trading hours. To avoid beginning or end of day affects and after hour auctions, we restrict ourselves to the limit orders placed between 9 and 4. We also restrict ourselves to non crossing orders (real limit orders).

We study over 2 years of transaction data translating into millions of transactions during the years 2000-2002 in the London Stock Exchange. We focus on four stocks, AZN, PRU, LLOY and VOD. These stocks provide a wide range of different characteristics including price, volume, spread and volatility.

Limit orders are measured as the distance from the current best. For buy orders, if the current best buy is Π_b and the buy limit order is π_b then the limit price is defined as $X = \pi_b - \Pi_b$ in terms of tick prices and $x = log(\pi_b/\Pi_b)$ for logarithms. Similarly for sell limit orders with a value of π_a and a current best of Π_a the values are defined as $X = \Pi_a - \pi_a$ and $x = log(\Pi_a/\pi_a)$ respectively. In both cases the spread is defined as $S = \Pi_a - \Pi_b$ in ticks and $s = log(\Pi_a/\Pi_b)$ in logarithms.

III. TICKS AND PERCENT

By looking at both ticks and logarithms, we can characterize whether traders "think" in terms of ticks or percentages. In other words, will the average variance of the distribution have a width or an exponent that is unaffected by the change in price. This would imply that

people think in ticks. On the other hand of the shape of the distribution as displayed in logarithmic values is more insensitive to the change in price, then limit order placers must be placing orders in the context of percentages.

In Figure 1, we plot the probability distribution of limit orders in log units $P(x)$ for the four stocks. This can be compared to Figure 2 in which we plot the probability distribution $P(X)$ in tick units. For both of these plots, it can be noted that the values are peaked at $x = 0$ and $X = 0$ respectively, showing that the most common bid is placed at the same best. Away from this peak, the distributions fall of rapidly with long tails.


FIG. 1: The probability distribution of unweighted limit orders placed in log units for several stocks. The colors correspond to AZN (black), LLOY (blue), PRU (green) and VOD (red).


FIG. 2: The probability distribution of unweighted limit orders placed in tick units for several stocks.The colors correspond to AZN (black), LLOY (blue), PRU (green) and VOD (red).

In the logarithms units, it appears that all distributions are fairly similar. Considering ticks, it appears that for limit orders placed within the book (outside of the spread) the distributions for different stocks are all very similar. Insider the spread, the distributions may vary. This is most likely due to the spread that has a different shape for each of these stocks. This idea is discussed more thoroughly in the next section. Finally, it is interesting to note that the distributions for buy and sell are very similar.

IV. WEIGHTING BY THE SPREAD

Limit orders placed within the spread are directly affected by spread distributions. We first review the probability distributions of the spread for several stocks in the following figures. In Figure 3 the probability distribution of the spread in logarithms $P(s)$ is plotted for the four stocks. This can be compared with Figure 4 in which the probability distribution of the spread in tick values $P(S)$ is plotted. The spread probability distribution changes for different stocks but the overall shape is similar. It consists of a peak at low values and a power law decay with an exponent of 3 for large values.


FIG. 3: The probability distribution of the spread in log units for several stocks.
The colors correspond to AZN (black), LLOY (blue), PRU (green) and VOD (red).

Order placement within the spread but close to the same best is where the spread distribution is high. In other words, it is easy to place limit orders in this region because the spread is usually bigger. However for large positive X, limit orders are more difficult to place because they are more likely to cross opposite best and become a market order.

To account for the inequality in how limit orders are placed within the spread, we employ the idea of weighting the distribution, P to P*. To model order placement in this range we look for a functional form $p(x|s)$, the probability of an event, x, in our case limit order placement, conditioned on the spread. We restrict ourselves only to orders, $x < s$, i.e., placed within the spread. For limit and market orders placed as $x >= s$ they are not considered part of the calculation. We can get bet-

FIG. 4: The probability distribution of the spread in tick units for several stocks. The colors correspond to AZN (black), LLOY (blue), PRU (green) and VOD (red).

ter statistical convergence by studying $p(x|s > 0)$. The sampling asymmetry can be dealt with as follows

$$p(x|s > s_1) = \frac{\int_0^\infty p(x|s)p(s)ds}{\int_0^\infty p(s)ds} \qquad (1)$$

There are effectively two ranges that must be considerd: $x < 0$ (limit orders placed inside the book) and $0 < x < s$ (orders placed inside the spread). For the first range, there is no weighting needed. For the second range, the weighted probability distribution can be written as

$$p(x) = p(x|s > s_1)\frac{\int_0^\infty p(s)ds}{\int_x^\infty p(s)ds} \qquad (2)$$

Using this definition, we can approximate $p(x)$ within the spread as

$$p(x) \approx p(x|s > 0)\frac{N(s > 0)}{N(s > x)} \qquad (3)$$

There are some subtleties about weighting the probability distribution in this matter that mainly come down to choosing bin widths of finite thickness. to account for this fact, we weight each limit order, x, by the weighting function separately. When grouping these orders into a bin, we choose the center of the bin to be the average value of the weighted limit orders within that bin.

For self consistency, we perform a similar weighting function on the spread in tick units. In this case, the weighting function is defined as

$$W(X) = \frac{N(S > S_1)}{N(S > X)} \qquad (4)$$

FIG. 5: The weighted probability distribution of limit orders in log units for all the stocks. The colors correspond to AZN (black), LLOY (blue), PRU (green) and VOD (red).


FIG. 6: A measure of the weighted probability distributions of positive x measured in tick. The colors correspond to AZN (black), LLOY (blue), PRU (green) and VOD (red).


FIG. 7: The weighted probability distribution of limit orders within the spread in tick units for AZN. The black is the unweighted probability distribution while the magenta is the weighted.

In Figure 5, we plot the weighted probability distribution in logarithms $P^*(x)$ for four different stocks. Is it possible that the distribution is nearly symmetric? If so why would this be the case? If this is possibly the case, what might be affecting the distribution from making it completely symmetric.

For comparison Figure 6 shows the weighted probability distribution for $X > 0$. In Figure 7 we plot the total probability distribution in ticks for both the weighted and unweighted case for AZN.

V. DISTRIBUTION SYMMETRY

In the previous section, we observed that some of the distributions appear to be nearly symmetric. We quantify this hypothesis by measuring $P(q)/P(-q)$ for a generalized variable, q. How close this measure stays to 1 is a measure of symmetry (What other approaches are there?

In Figure 8 we plot the ratio $P(X)/P(-X)$ for tick values in the weighted and unweighted case. As can be seen for the unweighted case this ratio is all well below one, while for the weighted case, this ratio is above one. This implies that there is no clear case for symmetry in the context of $P(X)$. For comparison, we plot the ratio $P(x)/P(-x)$ in Figure 9 for both the weighted and the unweighted cases. In this case, the distributions appear to be closer to symmetric, though some deviations are still apparent.

FIG. 8: A measure of the symmetry, $P(X)/P(-X)$ for both weighted and unweighted probability distributions measured in ticks. The colors correspond to AZN (black), LLOY (blue), PRU (green) and VOD (red).

VI. WHAT FACTORS AFFECT PLACEMENT

We conclude this discussion of limit order placement, by observing which market observables may affect placement. Although the distributions appear to be similar for

FIG. 9: A measure of the symmetry, $P(x)/P(-x)$ for both weighted and unweighted probability distributions measured in logarithms

FIG. 10: The probability distribution of limit order placement for three ranges of the spread: 0.00-0.002 (black), 0.002-0.005 (blue) and 0.005-0.010 (red). It appears that the shape of the distribution changes for small or high spreads, both above and below same best.

FIG. 11: The probability distribution of limit order placement for low (black) and high (red) volatility. More orders placed within the spread might imply that the spread is larger for high levels of volaitlity.

FIG. 12: The probability distribution of limit order placement for asks conditioned on whether the price moved up (red) or down (blue). More ask bids placed within the spread imply that price will go down. This suggests a self equilibriating, feedback mechanism in the market. The plots are switched for bids.

both buys and sells and for different stocks, etc. there may be other factors that affect placement.

In Figure 10, we plot the probability distribution of limit orders conditioned on the low, middle and high spread values. As can be seen, it does appear that spread distributions change limit order placement both above and below the same best. For low spread, limit orders placed below same best are more highly populated. In contrast, for large spread, there are far more orders placed within the spread and less placed deep in the book. Although only AZN sells is shown in the plot, we observed similar behavior for both buys as well as other stocks.

Figure 11 plots the probability for limit orders placed conditioned on volatility calculated over the past 10 minutes. The trend shown is similar to volatility calculated over longer times as well but is most notable on this time scale. For low volatility, there are more orders placed within the book and less in the spread. For high volatility, the converse is true. Again, this is also true for buys as well as for other stocks. It is possible that we are really observing the affect from spread fluctuations conditioned on volatility. In other words, when volatility is larger, the spread is more likely to be large. More work needs to be done to isolate this possibility.

Finally, we plot the placement of limit orders conditioned on price movement in Figure 12. When the price has gone up in the past ten minute, more ask orders are placed within the spread (implying that the price will go back down). On the other hand, when the price has recently gone down, more conservative orders are placed, providing more of a mechanism for the price to rise again. The opposite has been observed to be true for buy limit orders. Similar trends have been observed for the other stocks. This feature about limit order placement implies a self equilibriating feedback mechanism for order placement and market dynamics. Such trends imply that prices fluctuate around a constant value and markets are efficient.

VII. CONCLUSION

In this paper we have provided an extensive review of the statistical properties of limit order placements (orders which do not immediately transact) for four stocks on the London Stock Exchange. We look at limit order placements in the context of the distance from the current same best and explore the probability distributions in units of both ticks and logarithms (percentages). We review how orders placed within the spread need to be weighted by the probability that such a limit order could be placed and explore how this changes the probability distribution. We also look at a symmetry measure of weighted and unweighted distributions in both ticks and logarithms to explore the hypothesis that the placements are symmetric around same best. We conclude by reviewing the affect that other quantities such as volatility, price change and spread have on the placement of limit orders.

[1] J. Bouchaud, Mezard, and Potters, Quantitative Finance **2** (2002).
[2] J. Farmer and Zovko, Quantitative Finance **5**, 387 (2002).
[3] F. Lillo, Quantitative Finance **10**, 23 (2005).
[4] J. Farmer and Zovko, PNAS **102**, 287 (2005).