"Market making" behaviour in an electronic order book and its impact on the bid-ask spread

Ioane Muni Toke*

Chair of Quantitative Finance
Laboratory of Applied Mathematics and Systems
Ecole Centrale Paris, France

March 22, 2010

Abstract

It has been suggested that marked point processes might be good candidates for the modelization of financial high-frequency data. A special class of point processes, Hawkes processes, has been the subject of various investigations in the financial community. In this paper, we propose to enhance a basic order book simulator with limit and market orders arrival times following mutually (unsymmetrically) exciting Hawkes processes. Modelization is based on new empirical observations on interval times between orders that we verify on several markets (equity, bond futures, index futures). We show that this simple feature enables a much more realistic treatment of the bid-ask spread of the simulated order book.

Introduction

Orders' times of arrival : event time is not enough As of today, the study of arrival times of orders in an order book has not been a primary focus in order book modeling. Many toy models leave this dimension aside when trying to understand the complex dynamics of an order book. In most order driven market models such as [8, 18, 1], and in some order book models as well (e.g.[21]), a time step in the model is an arbitrary unit of time during with which many events may happen. We may call that clock *aggregated time*.

In most order book models such as [7, 19, 9], one order is simulated per time step with given probabilities. In the simple case where these probabilities are constant and independant of the state of the model, such a time treatment is equivalent to the assumption that order flows are homogeneous Poisson processes. [19] doesn't fall into this category since the probability to submit a

*Correspondance: Ecole Centrale Paris, Grande Voie des Vignes, 92290 Chatenay-Malabry, France. Email: ioane.muni-toke@ecp.fr

market order is not constant in time. In any case, these models use the clock known as *event time*.

A probable reason for this situation – leaving aside the fact that models can be sufficiently complicated without adding another dimension – is that many puzzling empirical observations (see e.g. [6] for a review of some of the well-known "stylized facts") can (also) be done in event time (e.g. autocorrelation of the signs of limit and market orders) or in aggregated time (e.g. volatility clustering).

However, it is clear that physical (calendar) time has to be taken into account for the modeling of a realistic order book model. For example, market activity is varies widely, and intraday seasonality is often observed as a well known U-shaped pattern. Even for a short time scale model – a few minutes, a few hours, since it may probably be illusory willing to build a model that would reproduce a whole trading day... – interarrival times of orders are very broadly distributed. The (Poisson) exponential distribution has to be discarded and these irregular flows of orders affect the empirical properties studied on order books.

Let us give one illustration. On figure 1, we plot examples of the empirical density function of the observed spread in event time (i.e. spread is measured each time an event happens in the order book), and in physical (calendar) time (i.e. measures are weighted by the time interval during which the order book is idle). It appears that density of the most probable values of the time-weighted

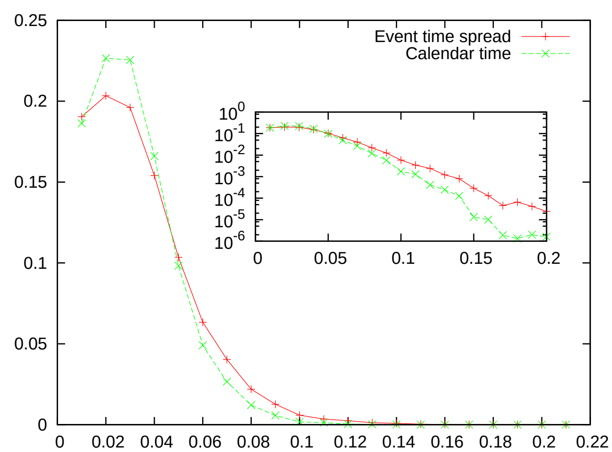

Figure 1: Empirical density function of the distribution of the bid-ask spread in event time and in calendar time. In inset, same data using a semi-log scale. This graph has been computed with 15 four-hour samples of tick data on the BNPP.PA stock (see section 1.1 for details).

distribution is higher than in the event time case. Symmetrically, the density of the least probable event is even smaller when time is taken into account. This tells us a few things about the dynamics of the order book, which could be summarized as follows : the wider the spread, the faster its tightening.

We can get another insight of this empirical property by measuring on our data the average waiting time before the next event conditionaly on the spread size. When computed on the lower one-third-quantile (small spread), the average waiting time is 320 milliseconds. When computed on the upper one-third-quantile (large spread), this average waiting time is 200 milliseconds. These observations complement some of the ones that can be found in the early paper [4].

Counting processes with dynamic intensity There is a trend in the econometrics literature advocating for the use of (marked) point processes for the modeling of financial timeseries. One may find a possible source of this interest in [11, 10], which introduce autoregressive conditional duration/intensity models (and test them on IBM trade times). [12] fit that type of models on the arrival times of limit, market and cancellation orders in an Australian stock market order book.

A particular class of point processes, known as the Hawkes processes, is of special interest for us, because of its simplicity of parametrization. A univariate linear self-exciting Hawkes process $(N_t)_{0<t<T}$, as introduced by [14, 15], is a point process with intensity :

$$\lambda(t) = \lambda_0 + \int_0^t \nu(t-s)dN_s, \tag{1}$$

where the kernel ν is usually parametrized as $\nu(t) = \alpha e^{-\beta t}$ or in a more general way $\nu(t) = \sum_{k=1}^{p} \alpha_k t^k e^{-\beta t}$. Statistics of this process are fairly well-known and results for a maximum likelihood estimation can be found in [20]. In a multivariate setting, mutual excitation is introduced. A bivariate model can thus be written:

$$\left\{ \begin{array}{lllll} \lambda^1(t) & = & \lambda_0^1 & + \int_0^t \nu_{11}(t-s)dN_s^1 & + \int_0^t \nu_{12}(t-s)dN_s^2 \\ \lambda^2(t) & = & \lambda_0^2 & + \int_0^t \nu_{21}(t-s)dN_s^1 & + \int_0^t \nu_{22}(t-s)dN_s^2 \end{array} \right. \tag{2}$$

The use of these processes in financial modeling is growing. We refer the reader to [3] for a review and [13] for a textbook treatment. In [5], a bivariate (generalized) Hawkes process is fitted to the timeseries of trades and mid-quotes events, using trading data of the General Motors stock traded on the New York stock Exchange. In [17] a ten-variate Hawkes process is fitted to the Barclay's order book on the London Stock Exchange, sorting orders according to their type and aggressiveness. It is found that the largest measured effect is the exciting effect of market orders on markets orders. [16] fits a bivariate Hawkes model to the buy and sell trades timeseries on the EUR/PLN (Euro/Polish Zlotych) FX market. Using the simplest parametrization of Hawkes processes and some (very) constraining assumptions, some analytical results of trade impact may be derived. [2] fits a bivariate Hawkes process to the trade timeseries of two different but highly correlated markets, the “Bund” and the “Bobl” (Eurex futures on mid- and long-term interest rates).

Organization of this paper In this paper, we propose to enhance a basic order book simulator with limit and market orders arrival times following mutually (unsymmetrically) exciting Hawkes processes. Modelization is based on new empirical observations, verified on several markets (equities, futures on index, futures on bonds), and detailed in section 1. More specifically, we observe evidence of some sort of market making in the studied order books: after a market order, a limit order is likely to be submitted more quickly than it would have been without the market order. In other words, there is a clear reaction that seems to happen once liquidity has been removed from the order book, a limit order is triggered to replace it. We also show that the reciprocal effect is not observed on the studied markets. These features lead to the use of unsymmetrical Hawkes processes for the modeling of a zero-intelligence agent based order book simulator described in section 2. We show in section 3 that this simple feature enables a much more realistic treatment of the bid-ask spread of the simulated order book.

1 Empirical evidence of “market making”

1.1 Data and observation setup

We use order book data for several types of financial assets :

- BNP Paribas (Reuters Identification Code : BNPP.PA) : 7th component of the CAC40 during the studied period
- Peugeot (Reuters Identification Code : PEUP.PA) : 38th component of the CAC40 during the studied period
- Lagardère SCA (Reuters Identification Code : LAGA.PA) : 33th component of the CAC40 during the studied period
- Dec.2009 futures on the 3-month Euribor (Reuters Identification Code : FEIZ9)
- Dec.2009 futures on the Footsie index (Reuters Identification Code : FFIZ9)

We use Reuters RDTH tick-by-tick data from September 10th, 2009 to September 30th, 2009 (i.e. 15 days of trading). For each trading day, we use only 4 hours of data, precisely from 9:30am to 1:30pm. As we are studying Europen markets, this time frame is convenient because it avoids the opening of American markets and the consequent increase of activity.

Our data is composed of snapshots of the first five limits of the order books (ten for the BNPP.PA stock). These snapshots are timestamped to the millisecond and taken at each change of any of the limits or at each transaction. For each snapshot, prices and volumes for all available limits and transactions are available. The data analysis is performed as follows for a given snapshot :

1. if the transaction fields are not empty, then we record a market order, with given price and volume;
2. if the quantity offered at a given price has increased, then we record a limit order at that price, with a volume equal to the difference of the quantities observed;
3. if the quantity offered at a given price has decreased without any transaction being recorded, then we record a cancellation order at that price, with a volume equal to the difference of the quantities observed;
4. finally, if two orders of the same type are recorded at the same time stamp, we record only one order with a volume equal to the sum of the two measured volumes.

Therefore, market orders are well observed since transactions are explicitly recorded, but it is important to note that our measure of the limit orders and cancellation orders is not direct. In table 1, we give for each studied order book the number of market and limit orders detected on our 15 4-hour samples. On

Code	Number of limit orders	Number of market orders
BNPP.PA	321,412	48,165
PEUP.PA	228,422	23,888
LAGA.PA	196,539	9,834
FEIZ9	110,300	10,401
FFIZ9	799,858	51,020

Table 1: Number of limit and markets orders recorded on 15 samples of four hours (Sep 10th to Sep 30th, 2009 ; 9:30am to 1:30pm) for 5 different assets (stocks, index futures, bond futures)

the studied period, market activity ranges from 2.7 trades per minute on the least liquid stock (LAGA.PA) to 14.2 trades per minute on the most traded asset (Footsie futures).

1.2 Empirical evidence of "market making"

Our idea for an enhanced model of order streams is based on the following observation: once a market order has been placed, the next limit order is likely to take place faster than usual. To illustrate this, we compute for all our studied assets:

- the empirical probability density function (pdf) of the inter arrival times of the counting process of all orders (limit orders and market orders mixed), i.e. the time step between any order book event (other than cancelation)

- and the empirical density function of the time step between a market order and the immediatly following limit order.

If an independant Poisson assumption held, then these empirical distribution should be identical. However, we observe a very high peak for the short inter-arrival times in the second case. The first moment of these empirical distributions is significant : one the studied assets, we find that the average time between a market order and the following limit order is 1.3 (BNPP.PA) to 2.6 (LAGA.PA) times shorter than the average time between two random consecutive events.

On the graphs shown in figure 2, we plot the full empirical distributions for four of the five studied assets[1]. We observe their broad distribution and the sharp peak for the shorter times : on the Footsie future market for example, 40% of the measured time steps between consecutive events are less that 50 milliseconds ; this figure jumps to nearly 70% when considering only market orders and their following limit orders. . .

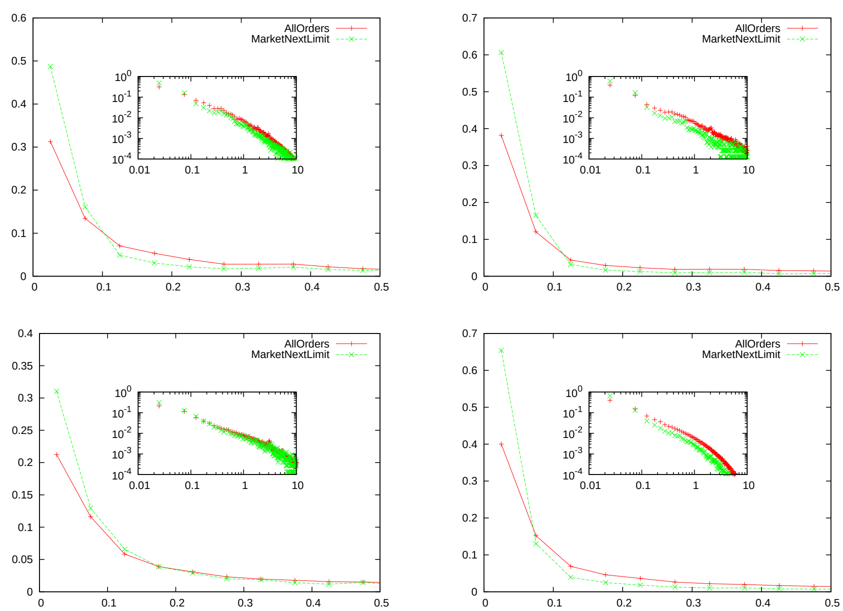

Figure 2: Empirical density function of the distribution of the timesteps between two consecutive orders (any type, market or limit) (full line) and empirical density function of the distribution of the time steps between a market order and the immediatly following limit order. X-axis is scaled in seconds. In insets, same data using a semi-log scale. Studied assets : BNPP.PA (top left), LAGA.PA (top right), FEIZ9 (bottom left), FFIZ9 (bottom right).

[1] Observations are identical on all the studied assets.

This observation is an evidence for some sort of market-making behaviour of some participants on those markets. It appears that the submission of market orders is monitored and triggers automatic limit orders that add volumes in the order book (and not far from the best quotes, since we only monitor the five best limits).

In order to confirm this finding, we perform non-parametric statistical test on the measured data. For all four studied markets, omnibus Kolmogorov-Smirnov and Cramer-von Mises tests performed on random samples establish that the considered distributions are statistically different. If assuming a common shape, a Wilcoxon-Mann-Withney U test clearly states that the distribution of interval times between a market orders and the following limit order is clearly shifted to the left compared to the distributions of interval times between any orders (see appendix B), i.e. the average "limit following market" reaction time is shorter than the average interval time between random consecutive orders.

Note also that these observations are still valid if we measure :

- the time steps between a (bid,ask) market order and the following limit orders on the same side of the order book (bid,ask);
- the time steps between a (bid,ask) market order and the following limit orders on the opposite side of the order book (ask, bid).

In other words, we cannot distinguish on the data if liquidity is added where the market order has been submitted or on the opposite side. On figure 3, we plot the empirical distributions of interval times between a market order and the following limit order on the same side or the following limit order on the oppposite side. It appears for all studied assets that both distributions are roughly identical.

Therefore, we do not infer any empirical property of placement: when a market order is submitted, the intensity of the limit order process increases *on both sides* of the order book.

1.3 A reciprocal "market following limit" effect ?

We now check if a similar or opposite distorsion is to be found on market orders when they follow limit orders. To investigate this, we compute for all our studied assets the "reciprocal" measures:

- the empirical probability density function (pdf) of the inter arrival times of the counting process of all orders (limit orders and market orders mixed), i.e. the time step between any order book event (other than cancelation)
- and the empirical density function of the time step between a market order and the previous limit order.

As previously, if an independant Poisson assumption held, then these empirical distribution should be identical. Results for four assets are shown on figure 4. Contrary to previous case, no effect is very easily interpreted. For the three

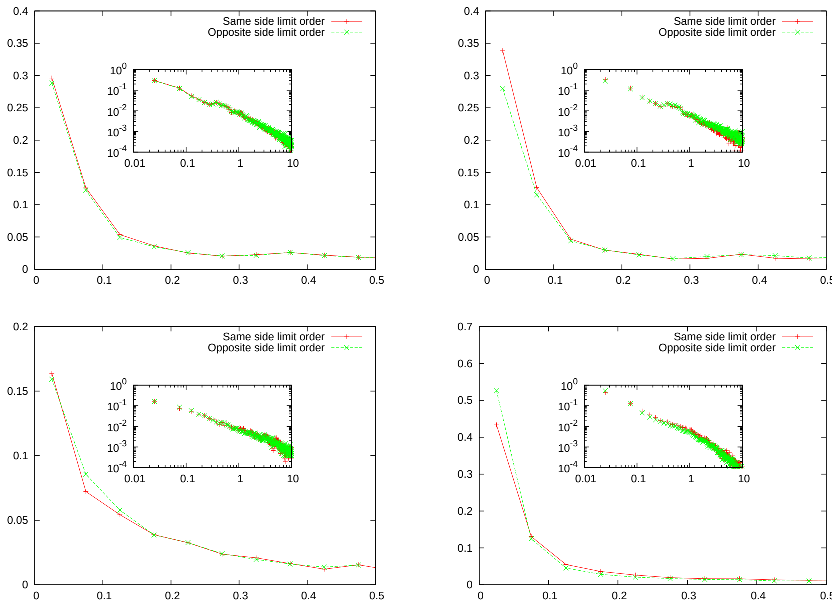

Figure 3: Empirical density function of the distribution of the time steps between a market order and the immediatly following limit order on the same side and on the opposite side. X-axis is scaled in seconds. In insets, same data using a semi-log scale. Studied assets : BNPP.PA (top left), PEUP.PA (top right), FEIZ9 (bottom left), FFIZ9 (bottom right).

stocks (BNPP.PA, LAGA.PA and PEUP.PA (not shown)), it seems that the empirical distribution is less peaked for small interval times, but difference is much less important than in the previous case. As for the FEI and FFI markets, the two distributions are even much closer. Non-parametric tests confirms these observations.

Performed on data from the three equity markets, Kolmogorov tests indicate different distribution and Wilcoxon tests enforce the observation that interval times between a limit order and a following market order is stochastically larger than the interval time between unidentified orders. As for the future markets on Footsie (FFI) and 3-month Euribor (FEI), Kolmogorov tests does not indicate differences in the two observed distributions, and the result is confirmed by a Wilcoxon test that concludes at the equality of the means. Results are given in appendix B.

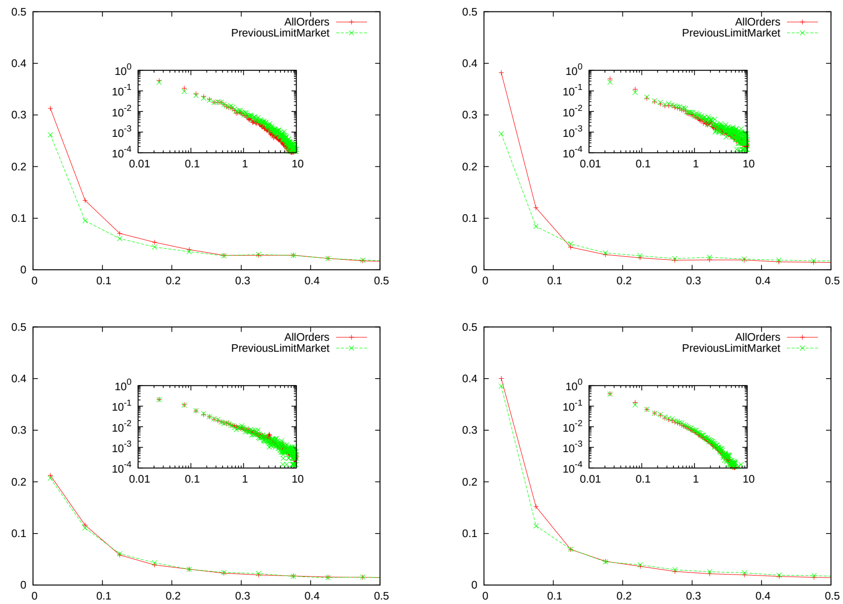

Figure 4: Empirical density function of the distribution of the timesteps between two consecutive orders (any type, market or limit) (full line) and empirical density function of the distribution of the time steps between a limit order and an immediatly following market order. In insets, same data using a semi-log scale. Studied assets : BNPP.PA (top left), LAGA.PA (top right), FEIZ9 (bottom left), FFIZ9 (bottom right).

1.4 Addressing potential high-frequency artifacts : zooming out

One might be concerned that this kind of empirical study relies too heavily on the accuracy of high-frequency data, accuracy which might be questionable for very liquid stocks. Anyone who ever studied high-frequency trading data knows that reconstructing a trading sequence is a difficult process that uncovers many recording errors.

In particular, one might be suspicious that the effects described in sections 1.2 and 1.3, despite all the care taken in cautiously checking the orders, might be artifacts due to flaws in the data recording process (e.g. trades and quotes asynchronicity) or possible flaws in the order identification process (e.g. falsely detected quotes updates after trades).

In order to address theses concerns, we propose here another statistical evidence of the phenomena, but obtained at a larger time scale. We split each of our 15 samples of four-hour trading data in N intervals of length $\Delta T = 1$ second. For each interval $n = 0, \ldots, N$, we count the number M_n of market orders and the number L_n of limit orders. We then compute, conditionally on the number of market orders in a given interval, the average number of limit orders in this interval and in adjacent intervals, i.e. $E[L_n|M_n]$, $E[L_{n-1}|M_n]$ and $E[L_{n+1}|M_n]$. Results are shown in figure 5. Of course we should observe that

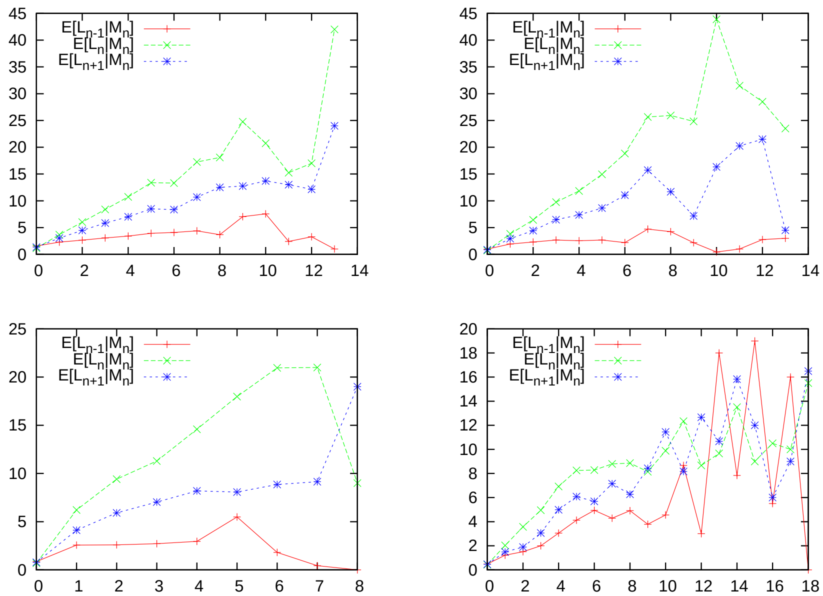

Figure 5: Expectation of the number of limit orders in adjacent intervals conditionally on the number of market orders counted in a given interval. Studied assets : BNPP.PA (top left), PEUP.PA (top right), LAGA.PA (bottom left), FEIZ9 (bottom right).

the number of limit orders in a given interval and also in the previous one and next one is more or less linked with the number of market orders counted in this interval. This is simply due to the intraday seasonality of activity, and the clustered patterns of arrival of orders. And as expected, it appears clearly that the number of limit and market orders in a given interval are highly correlated: the more market orders are counted, the more limit orders are counted in the same interval (green dash-x line on figure 5).

But the intensity of this link is not the same when looking at the previous or next intervals. In each case, the growth is roughly linear (data with large number of market order are the noisiest). By performing a simple linear regression

$$E[L_{n+i}|M_n] = \alpha_i M_n + \beta_i, i \in \{-1, 0, +1\},$$

we can illustrate these differences. For example for the BNPP.PA stock, we obtain the following coefficients: $\alpha_{-1} = 0.47, \alpha_0 = 2.3, \alpha_{+1} = 1.4$.

It is thus clear that the number of market orders in a given interval has an influence on the number of limit orders on the next one: the more market orders in a given interval, the more limit orders in the following interval (blue dash-star line). But it is only weakly linked, if a link is observed at all, to the number of limit orders observed in the previous interval (red full line on figure 5). As in the previous sections, these empirical facts are clear for equity stocks, and still observed but less strikingly for future assets.

Furthermore, we can also verify the results of section 1.3 and check that even at a larger time scale there is no reciprocal effect "limit orders triggering market orders". On figure 6, we plot the empirical conditional means $E[M_{n-1}|L_n]$, $E[M_n|L_n]$ and $E[M_{n+1}|L_n]$. It is clear that a large number of limit orders in

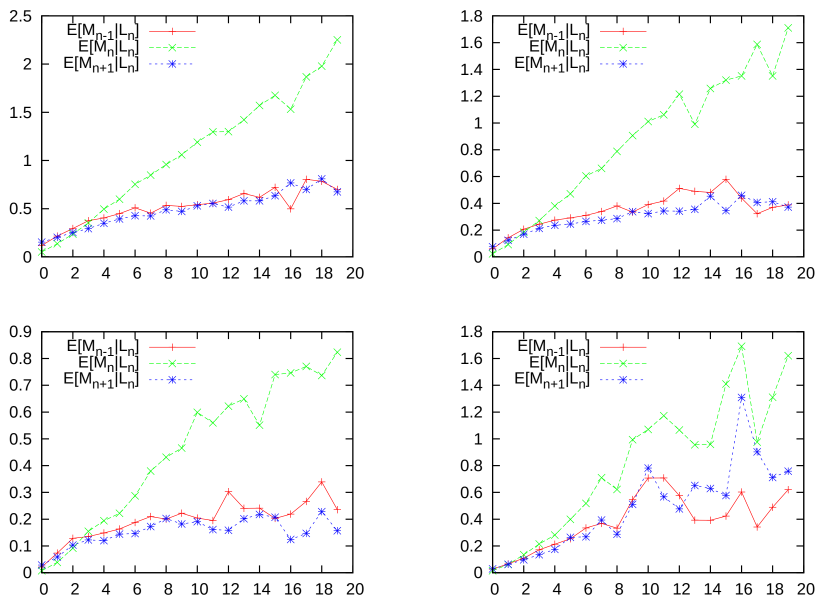

Figure 6: Empirical mean of the number of market orders in adjacent intervals conditionally on the number of limit orders counted in a given interval. Studied assets : BNPP.PA (top left), PEUP.PA (top right), LAGA.PA (bottom left), FEIZ9 (bottom right).

a given interval is only weakly linked to the number of market orders in the previous or in the following interval. With these measures, we cannot even distinguish between the previous and the following interval.

We thus have shown that the "market making" or "liquidity replenishment" effect is still observed at a larger time scale. Using interval of one-second length (instead of tens of milliseconds), we still find that the intensity of the limit order process increases with the number of market orders counted, and that the reciprocal effect is not observed.

***

Let us summarize the empirical findings. We observe that the interval time between a market order and the following limit order is stochastically much shorter than the interval time between two undistinguished orders (see section 1.2) and that the reciprocal statement is not true : the distribution of the interval times between a market order and the previous limit order is not clearly different from the interval time between two unidentified orders. Theses empirical facts can still be observed at a larger time scale where high-frequency artifacts can be ruled out. These observations enforce the key role of the market orders that really seem to drive the market activity.

2 Order book simulators with mutually exciting order flows

Following these previous observations, we enhance a basic agent-based order book simulator with dependance between the flows of limit and market orders.

2.1 The basic Poisson model

We use as base model a standard zero-intelligence agent-based market simulator built as follows. One agent is a liquidity provider. This agent submits limit orders in the order books, orders which he can cancel at any time. This agent is simply characterized by a few random distributions :

1. submission times of new limit orders are distributed according to a homogeneous Poisson process N^L with intensity λ^L;

2. submission times of cancelation of orders are distributed according to homogeneous Poisson process N^C with intensity λ^C;

3. new limit orders' placement around from the same side best quote follows a Student's distribution with degrees of freedom ν_1^P, shift parameter m_1^P and scale parameter s_1^P;

4. new limit orders' volume is randomly distributed according to an exponential law with mean m_1^V;

5. in case of a cancelation, the agent deletes his own orders with probability δ.

The second agent in the basic model is a noise trader. This agent only submits market order (it is sometimes refered to as the liquidity taker). Characterization of this agent is even simpler:

6. process of market orders' submission times is a homogeneous Poisson process N^M with intensity μ;

7. market orders' volume is randomly distributed according to an exponential law with mean m_2^V.

For all the experiments, agents submit orders on the bid or the ask side with probability 0.5. This basic model will be henceforth refered to as "HP" (Homogeneous Poisson).

Assumptions 1, 2 and 6 (Poisson) will be replaced in our enhanced model. Assumption 3 (Student) is in line with empirical observations in [19]. Assumptions 4 and 7 are in line with empirical observations in [6] as far as the main body of the distribution is concerned, but fail to represent the broad distribution observed in empirical studies. All the parameters except δ, which we kept exogenous, can be more or less roughly estimated on our data. In fact δ is the parameter of the less realistic feature of this simple model, and is thus difficult to calibrate. It can be used as a free parameter to fit the realized volatility.

2.2 Adding dependance between order flows

We have found in section 1.2 that market data shows that the flow of limit orders strongly depends on the flow of market order. We thus propose that in our experiment, the flow of limit and market orders are modeled by Hawkes processes N^L and N^M, with stochastic intensities respectively λ and μ defined as:

$$\left\{ \begin{array}{rclll} \mu^M(t) & = & \mu_0^M & + \int_0^t \alpha_{MM} e^{-\beta_{MM}(t-s)} dN_s^M & \\ \lambda^L(t) & = & \lambda_0^L & + \int_0^t \alpha_{LM} e^{-\beta_{LM}(t-s)} dN_s^M & + \int_0^t \alpha_{LL} e^{-\beta_{LL}(t-s)} dN_s^L \end{array} \right. \tag{3}$$

Three mechanisms can be used here. The first two are self-exciting ones, "MM" and "LL". They are a way to translate into the model the observed clustering of arrival of market and limit orders and the broad distributions of their inter arrival times. In the empirical study [17], it is found that the measured excitation MM is important. In our simulated model, we will show (see 3.2) that this allows a simulator to provide realistic distributions of trades' interarrival times.

The third mechanism, "LM", is the direct translation of the empirical property we've presented in section 1.2. When a market order is submitted, the intensity of the limit order process N^L increases, enforcing the probability that a "market making" behaviour will be the next event. We do no implement the reciprocal mutual excitation "ML", since we do not observe that kind of influence on our data as explained in section 1.3.

Rest of the model is unchanged. Turning these features successively on and off gives us several models to test – namely HP (Homogeneous Poisson processes), LM, MM, MM+LM, MM+LL, MM+LL+LM – to try to understand the influence of each effect.

3 Numerical results on the order book

3.1 Fitting and simulation

We fit these processes by computing the maximum likelihood estimators of the parameters of the different models on our data. Results are given in Appendix A. As expected, estimated values varies with the market activity on the day of the sample. However, it appears that estimation of the parameters of stochastic intensity for the MM and LM effect are quite robust. We find an average relaxation parameter $\hat{\beta}_{MM} = 6$, i.e. roughly 170 milliseconds as a characteristic time for the MM effect, and $\hat{\beta}_{LM} = 1.8$, i.e. roughly 550 milliseconds characteristic time for the LM effect. Estimation of models including the LL effect are much more troublesome on our data. In the simulations that follows, we assume if needed that the self-exciting parameters are similar ($\alpha_{MM} = \alpha_{LL}$, $\beta_{MM} = \beta_{LL}$) and ensure that the number of market orders and limit orders in the different simulations is roughly equivalent (i.e. approximately 145000 limit orders and 19000 market orders for 24 hours of continuous trading). Table 2 summarizes the numerical values used for simulation. Fitted parameters are in agreement with an assumption of asymptotical stationarity.

Model	μ_0	α_{MM}	β_{MM}	λ_0	α_{LM}	β_{LM}	α_{LL}	β_{LL}
HP	0.22	-	-	1.69	-	-	-	-
LM	0.22	-	-	0.79	5.8	1.8	-	-
MM	0.09	1.7	6.0	1.69	-	-	-	-
MM LL	0.09	1.7	6.0	0.60	-	-	1.7	6.0
MM LM	0.12	1.7	6.0	0.82	5.8	1.8	-	-
MM LL LM	0.12	1.7	5.8	0.02	5.8	1.8	1.7	6.0

Common parameters :	$m_1^P = 2.7, \nu_1^P = 2.0, s_1^P = 0.9$
	${}_1^V = 275, m_2^V = 380$
	$\lambda^C = 1.35, \delta = 0.015$

Table 2: Estimated values of parameters used for simulations.

We compute long runs of simulations with our enhanced model, simulating each time 24 hours of continuous trading. Note that using the chosen parameters, we never face the case of an empty order book. We observe several statistics on the results, which we discuss in the following sections.

3.2 Impact on arrival times

We can easily check that introducing self- and mutually exciting processes into the order book simulator helps producing more realistic arrival times. Figure 7 shows the distributions of inter arrival times of market orders (left) and limit orders (right). As expected, we check that the Poisson assumption has to be discarded, while the Hawkes processes help getting more weight for very short inter arrival times.

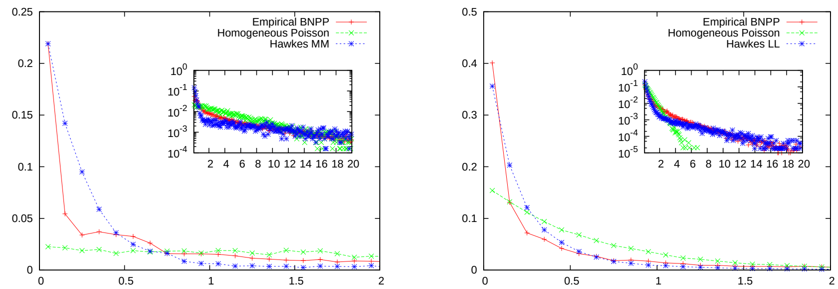

Figure 7: Empirical density function of the distribution of the interarrival times of market orders (left) and limit orders (right) for three simulations, namely HP, MM, LL, compared to empirical measures. In inset, same data using a semi-log scale. Numerical details are given in section 1.1 for the empirical distribution, and in section 3.1 for the simulated ones.

We also verify that models with only self-exciting processes MM and LL are not able to reproduce the "market making" feature described in section 1.2. Distribution of interval times between a market order and the next limit order are plotted on figure 8. As expected, no peak for short times is observed if the LM effect is not in the model. But when the LM effect is included, the simulated distribution of interval times between a market order and the following limit order is very close to the empirical one.

3.3 Impact on the bid-ask spread

Besides a better simulation of the arrival times of orders, we argue that the LM effect also helps simulating a more realistic behavior of the bid-ask spread of the order book. On figure 9, we compare the distribution of the spread for three models – HP, MM, MM+LM – in regard to the empirical measures. We first observe that the model with homogeneous Poisson processes produces a fairly good shape for the spread distribution, but slightly shifted to the right. Small spread values are largely underestimated. When adding the MM effect in order to get a better grasp at market orders' arrival times, it appears that we flatten the spread distribution. One interpretation could be that when the process N^M is excited, markets orders tend to arrive in cluster and to hit the first limits of the order book, widening the spread and thus giving more weight to large spread values. But since the number of orders is roughly constant in our simulations, there has to be periods of lesser market activity where limit orders reduce the spread. Hence a flatter distribution.

Here, the MM+LM model produces a spread distribution much closer to the empirical shape. It appears from figure 9 that the LM effect reduces the spread: the "market making" behaviour, i.e. limit orders triggered by market orders, helps giving less weight to larger spread values (see the tail of the distribution)

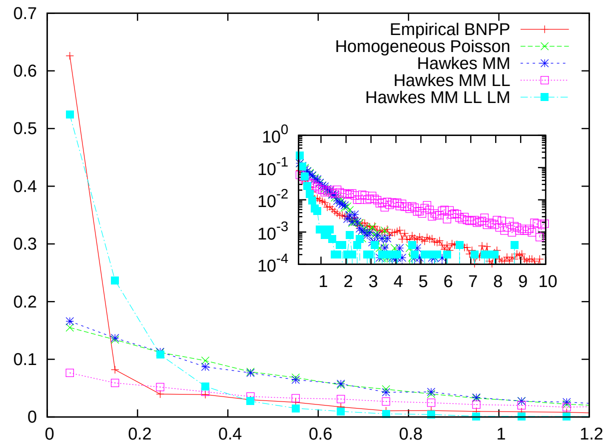

Figure 8: Empirical density function of the distribution of the interval times between a market order and the following limit order for three simulations, namely HP, MM+LL, MM+LL+LM, compared to empirical measures. In inset, same data using a semi-log scale. Numerical details are given in section 1.1 for the empirical distribution, and in section 3.1 for the simulated ones.

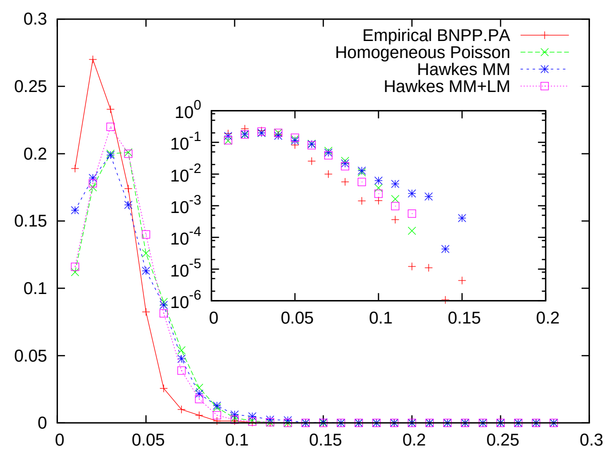

Figure 9: Empirical density function of the distribution of the bid-ask spread for three simulations, namely HP, MM, MM+LM, compared to empirical measures. In inset, same data using a semi-log scale. X-axis is scaled in euro (1 tick is 0.01 euro). Numerical details are given in section 1.1 for the empirical distribution, and in section 3.1 for the simulated ones.

and to sharpen the peak of the distribution for small spread values. Thus, it seems that simulations confirm the empirical properties of a "market making" behaviour on electronic order books.

We show on figure 10 that the same effect is observed in an even clearer way with the MM+LL and MM+LL+LM models. Actually, the spread dis-

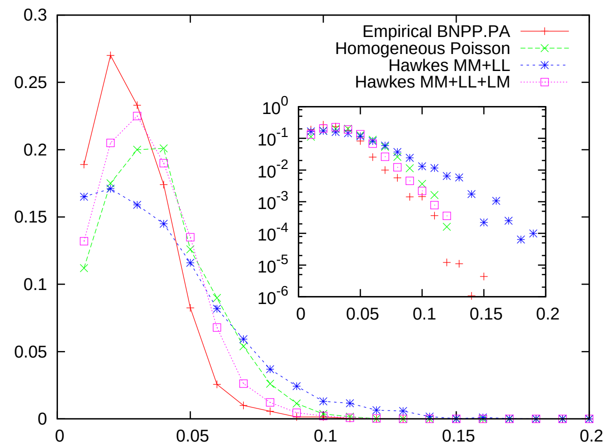

Figure 10: Empirical density function of the distribution of the bid-ask spread three simulations, namely HP, MM, MM+LM, compared to empirical measures. In inset, same data using a semi-log scale. X-axis is scaled in euro (1 tick is 0.01 euro). Numerical details are given in section 1.1 for the empirical distribution, and in section 3.1 for the simulated ones.

tribution produced by the MM+LL model is the flattest one. This is in line with our previous argument. When using two independant self exciting Hawkes processes for arrival of orders, moments of high market orders' intensity gives more weight to large spread values, moments of high limit orders' intensity gives more weight to small spread values. Adding the cross-term LM to the processes implements a coupling effect that helps reproducing the empirical shape of the spread distribution. The MM+LL+LM simulated spread is the closest to the empirical one.

3.4 Limitations of the model

It is somewhat remarkable to observe that these variations of the spread distributions is done with little or no change in the distributions of the variations of the mid-price. As shown on figure 11, the distributions of the variations of the mid-price sampled every 30 seconds are nearly identical for all the simulated models (and much tighter than the empirical one). This is due to the fact that the simulated order books are much more furnished than the empirical one, hence the smaller standard deviation of the mid price variations. One solution

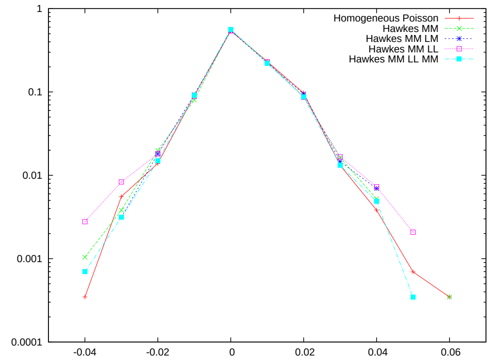

Figure 11: Empirical density function of the distribution of for the 30-second variations of the mid-price for five simulations, namely HP, MM, MM+LM, MM+LL, MM+LL+LM. In inset, same data using a semi-log scale. X-axis is scaled in euro (1 tick is 0.01 euro). Numerical details are given in section 1.1 for the empirical distribution, and in section 3.1 for the simulated ones.

to get thinner order books and hence more realistic value of the variations of the mid-price would be to increase our exogenous parameter δ. But in that case, mechanisms for the replenishment of an empty order book should be carefully studied, which is still to be done.

4 Conclusion and future work

We have shown the the use of Hawkes processes may help producing a realistic shape of the spread distribution in a zero-intelligence agent-based order book simulator. We emphasize on the role of the excitation of the limit order process by the market order process. This coupling of the processes, similar to a "market making" behaviour, is empirically observed on several markets, and simulations confirms it is a key component for realistic order book models.

Future work should investigate if other processes or other kernels (ν_{LM} in our notation) might better fit the observed orders flows. In particular, we observe very short characteristic times, which should lead us to question the use of the exponential decay. Furthermore, as pointed out in the paper, many other mechanisms are to be investigated : excitation of markets orders, link with volumes, replenishment of an empty order book, etc.

Acknowledgement The author is indebted to Frédéric Abergel for initiating this work and for many insights and discussions, and to Emmanuel Bacry and Charles-Albert Lehalle for very helpful comments.

References

[1] V. Alfi, M. Cristelli, L. Pietronero, and A. Zaccaria, *Minimal agent based model for financial markets i - origin and self-organization of stylized facts*, The European Physical Journal B **67** (2009), no. 3, 13 pages.

[2] Emmanuel Bacry, *?*, Finance and Statistics (Fiesta) Seminar, Ecole Polytechnique, 2010.

[3] Luc Bauwens and Nikolaus Hautsch, *Modelling financial high frequency data using point processes*, Handbook of Financial Time Series, 2009, pp. 953–979.

[4] Bruno Biais, Pierre Hillion, and Chester Spatt, *An empirical analysis of the limit order book and the order flow in the paris bourse*, The Journal of Finance **50** (1995), no. 5, 1655–1689.

[5] Clive G. Bowsher, *Modelling security market events in continuous time: Intensity based, multivariate point process models*, Journal of Econometrics **141** (2007), no. 2, 876–912.

[6] Anirban Chakraborti, Ioane Muni Toke, Marco Patriarca, and Frederic Abergel, *Econophysics: Empirical facts and agent-based models*, 0909.1974 (2009).

[7] D. Challet and R. Stinchcombe, *Analyzing and modeling 1+ 1d markets*, Physica A: Statistical Mechanics and its Applications **300** (2001), no. 1-2, 285–299.

[8] Rama Cont and Jean-Philipe Bouchaud, *Herd behavior and aggregate fluctuations in financial markets*, Macroeconomic Dynamics **4** (2000), no. 02, 170–196.

[9] Rama Cont, Sasha Stoikov, and Rishi Talreja, *A Stochastic Model for Order Book Dynamics*, SSRN eLibrary (2008) (English).

[10] Robert F. Engle, *The econometrics of Ultra-High-Frequency data*, Econometrica **68** (2000), no. 1, 1–22.

[11] Robert F. Engle and Jeffrey R. Russell, *Forecasting the frequency of changes in quoted foreign exchange prices with the autoregressive conditional duration model*, Journal of Empirical Finance **4** (1997), no. 2-3, 187–212.

[12] Anthony D. Hall and Nikolaus Hautsch, *Modelling the buy and sell intensity in a limit order book market*, Journal of Financial Markets **10** (2007), no. 3, 249–286.

[13] Nikolaus Hautsch, *Modelling irregularly spaced financial data*, Springer, 2004.

[14] Alan G. Hawkes, *Spectra of some self-exciting and mutually exciting point processes*, Biometrika **58** (1971), no. 1, 83–90.

[15] Alan G. Hawkes and David Oakes, *A cluster process representation of a Self-Exciting process*, Journal of Applied Probability **11** (1974), no. 3, 493–503.

[16] P. Hewlett, *Clustering of order arrivals, price impact and trade path optimisation*, Workshop on Financial Modeling with Jump processes, Ecole Polytechnique, 2006.

[17] Jeremy Large, *Measuring the resiliency of an electronic limit order book*, Journal of Financial Markets **10** (2007), no. 1, 1–25.

[18] T. Lux, M. Marchesi, and G. Italy, *Volatility clustering in financial markets*, Int. J. Theo. Appl. Finance **3** (2000), 675–702.

[19] S. Mike and J. D. Farmer, *An empirical behavioral model of liquidity and volatility*, Journal of Economic Dynamics and Control **32** (2008), no. 1, 200–234.

[20] Y. Ogata, *On Lewis' simulation method for point processes*, IEEE Transactions on Information Theory **27** (1981), no. 1, 23–31.

[21] T. Preis, S. Golke, W. Paul, and J. J. Schneider, *Multi-agent-based order book model of financial markets*, EPL (Europhysics Letters) **75** (2006), no. 3, 510–516.

A Fitting results

Date	μ	λ_0	α_{LM}	β_{LM}
Date	mu	lambda0	alpha	beta
20090910	0.27	0.69	3.85	1.82
20090911	0.19	0.42	3.25	1.43
20090914	0.18	0.57	2.85	1.1
20090915	0.19	0.58	4.33	1.96
20090916	0.31	0.56	4.61	1.97
20090917	0.27	0.61	4.15	1.5
20090918	0.28	0.91	9.14	2.88
20090921	0.19	1.0	7.59	1.58
20090922	0.14	0.8	6.41	1.27
20090923	0.15	0.86	7.76	1.92
20090924	0.13	1.0	8.45	1.44
20090925	0.15	0.96	6.88	1.27
20090928	0.16	0.98	5.53	1.22
20090929	0.38	0.89	6.31	2.91
20090930	0.35	1.05	6.99	2.9
	0.22	0.79	5.87	1.81

Table 3: Estimators for the LM model computed on the 15 4-hour sample of BNPP.PA stock data. Last line computes the average values of these parameters.

B Non-parametric tests on interval times distributions

	KS		WMWU	
Asset	D	p-value	U	1%-interval
BNPP	5.08	0	-9.59	[-2.58;2.58]
PEUP	6.05	0	-9.94	[-2.58;2.58]
LAGA	6.11	0	-11.21	[-2.58;2.58]
FEI	5.99	0	-7.5	[-2.58;2.58]
FFI	4.20	0	-14.3	[-2.58;2.58]

Table 4: Statistics computed for Kolmogorov-Smirnov (KS) test and Wilcoxon-Mann-Whitney U (WMWU) test for the empirical distributions of interval times between unidentified orders and interval times between a market order and the following limit order (see details in section 1.2. KS test is performed on all data binned by steps of size 0.05. MWMU is performed on random samples of size 10^3.

	KS		WMWU	
Asset	D	p-value	U	1%-interval
BNPP	2.90	4E-8	4.57	[-2.58;2.58]
PEUP	2.98	1E-7	4.95	[-2.58;2.58]
LAGA	5.32	0	5.72	[-2.58;2.58]
FEI	1.07	0.20	0.37	[-2.58;2.58]
FFI	0.86	0.43	-0.95	[-2.58;2.58]

Table 5: Statistics computed for Kolmogorov-Smirnov (KS) test and Wilcoxon-Mann-Whitney U (WMWU) test for the empirical distributions of interval times between unidentified orders and interval times between a limit order and the following market order (see details in section 1.3. KS test is performed on all data binned by steps of size 0.05. MWMU is performed on random samples of size 10^3.