

Can a Zero-Intelligence Plus Model Explain the Stylized Facts of Financial Time Series Data?

Imon Palit
SEFeMEQ Department
University of Rome
Tor Vergata
palit@uniroma2.it

Steve Phelps
Centre for Computational
Finance and Economic Agents
(CCFEA)
sphelps@essex.ac.uk

Wing Lon Ng
Centre for Computational
Finance and Economic Agents
(CCFEA)
wlng@essex.ac.uk

ABSTRACT

Many agent-based models of financial markets have been able to reproduce certain stylized facts that are observed in actual empirical time series data by using “zero-intelligence” agents whose behaviour is largely random in order to ascertain whether certain phenomena arise from market micro-structure as opposed to strategic behaviour. Although these models have been highly successful, it is not surprising that they are unable to explain *every* stylized fact, and indeed it seems plausible that although some phenomena arise purely from market micro-structure, other phenomena arise from the behaviour of the participating agents, as suggested by more complex agent-based models which use agents endowed with various forms of strategic behaviour. Given that both zero-intelligence and strategic models are each able to explain various phenomena, an interesting question is whether there are hybrid, “zero-intelligence plus” models containing a minimal amount of strategic behaviour that are simultaneously able to explain all of the stylized facts. We conjecture that as we gradually increase the level of strategic behaviour in a zero-intelligence model of a financial market we will obtain an increasingly good fit with the stylized facts of empirical financial time-series data. We test this hypothesis by systematically evaluating several different experimental treatments in which we incrementally add minimalist levels of strategic behaviour to our model, and test the resulting time series of price returns for the following statistical features: fat tails, volatility clustering, persistence and non-Gaussianity. Surprisingly, the resulting “zero-intelligence plus” models do *not* introduce more realism to the time series, thus supporting other research which conjectures that some phenomena in the financial markets are indeed the result of more sophisticated learning, interaction and adaptation.

Categories and Subject Descriptors

H.4 [Information Systems Applications]: Miscellaneous

1. INTRODUCTION

Many agent-based models of financial markets have made use of zero-intelligence agents in order to explain various

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

Copyright © 2012, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

phenomena in microeconomics and finance. In these models, market scenarios are simulated with agents whose behaviour is largely random. They are particularly useful in attempting to attribute whether a particular phenomenon is caused by the market mechanism or to the strategic behaviour of participating traders, since any non-random regularities that appear in the macroscopic behaviour of such a model can be attributed to the mechanics of the underlying market microstructure rather than to the behaviour of the agents.

Zero-intelligence models have been applied to studying phenomena in both economics and finance. The concept was originally introduced in the field of microeconomics in order to explain how traders were able to converge on equilibrium prices in a continuous double-auction trading environment: Gode and Sunder set out to determine whether this phenomena could be ascribed to the intelligence of the human traders, or alternatively whether the efficiency of the market was due to the trading institution itself [7]. They did so by introducing one of the first agent-based simulation models in which agents with various levels of *random* price setting behaviour were simulated under the same market microstructure rules as used in the original experiment with human traders. They found that they were able to reproduce convergence to equilibrium prices (under certain treatments), suggesting that the efficiency of the market was due to the “intelligence” of the market institution itself, rather than intelligent behaviour on the part of the participants.

However, in later work Cliff and Bruten [4] showed that under different treatments Gode and Sunder’s zero-intelligence agents were not able reproduce the original convergence results (“zero is not enough”), and that so called zero-intelligence *plus* (ZIP) agents with a minimal level of *conditional* bidding behaviour were required in order to converge to equilibrium prices under a wider variety of initial conditions.

Although zero-intelligence models originated in the field of microeconomics, there have been considerable successes in using them to explain the stylized facts of *financial* time series data. For example, one such stylized fact is positive correlation in order flow: the probability of observing a given type of order in the future is positively correlated with its empirical frequency in the past. This phenomena was originally documented by Bias et al. on the Paris Bourse exchange [1]. [8] was able to reproduce this stylized fact using a zero-intelligence agent-based model, thus suggesting that empirical regularities in order-flow arise from the operation of the market mechanism rather than the behaviour

of traders.

However, it would be naive to suppose that *all* phenomena observed in financial time series data could be reproduced using models in which traders behave entirely randomly, and indeed it seems apriori reasonable that although some phenomena arise purely from market micro-structure, other phenomena arise from the behaviour of the participating agents, as suggested by more complex agent-based models which use agents endowed with various forms of strategic behaviour [3].

An interesting question then arises as to what is the minimal level of intelligence required for an agent-based model to reproduce statistically-realistic time series data. Analogous to Cliff’s “zero is not enough” conjecture, we hypothesise that by gradually increasing the level of strategic behaviour in a zero-intelligence model we will obtain an increasingly good fit with the empirical data of financial markets. We test this hypothesis by systematically evaluating several different experimental treatments in which we incrementally add minimal amounts of conditional behaviour to our model, and test the resulting time series of returns for the following stylized facts: fat tails, volatility clustering, high-frequency persistence and non-Gaussianity [5][12].

Stylized facts are statistical regularities that are often found in real market data across different markets and different periods of time [5]. By reproducing these empirically observed stylized facts we can validate our model and by building a microscopic (trader-based) model of the market attempt to understand what assumptions are consistent with these properties. We test whether minimal strategic behaviour is sufficient to reproduce non-Gaussian distributions (characterised by excess kurtosis), fat tailed distributions (characterised by power law decay in the tail of the distribution), high frequency persistence (characterised by super-diffusive behaviour at short time scales) and volatility clustering (characterised by nonstationarity in price changes).

We start with a simple zero-intelligence model in which different *types* of order, as classified by the original Paris Bourse study [1] are submitted to the exchange. In our first experimental treatment order types are chosen from a discrete uniform distribution. We then gradually introduce additional complexity into the model by allowing the probability with which an event type is chosen to change in response to the state of the market – that is, we introduce *conditional* or strategic behaviour into our model. We do so systematically and in line with empirically-observed phenomena in real financial markets.

The paper is structured as follows. In sections 2 and 3 we give an overview of empirically documented conditional order flow phenomena in financial time series data. In section 4 we give detailed description of our model, including the event types and detailed descriptions of the seven experimental treatments used in our investigation. In section 5 we describe our methodology. In sections 6 and 7 we show and discuss our results, and finally we conclude in section 8.

2. ABSOLUTE AND CONDITIONAL ORDER FLOW

[1] show that there is a “diagonal effect” in high-frequency financial time-series data: viz., order flow is conditional on past order flow. Both [1] and [2] show there is a “stimulated refill” liquidity process: order flow is conditional on the state

of the book. Previous analysis by [13] equating effective costs of market and limit orders, demonstrated that a necessary condition for statistical arbitrage efficiency is a linear relationship between impact, bid-ask spread and volatility. [13] empirically found this relationship to hold true and propose this as a market law or stylized fact that has a theoretical underpinning based on ecology between liquidity takers and providers in pure order driven markets. They argue this phenomenon emerges as liquidity in an order-driven market self-organises toward statistical efficiency. Interestingly this theoretical motivation occurs endogenously without any role of information and does not espouse the traditional view of efficient market theory: that of informationally-efficient markets.

Using our model we investigate alternative treatments corresponding to both conditional and unconditional order flow. In each case, we test for fat tails, volatility clustering, long-memory and non-Gaussianity in returns. We do this by simulating the effects of adding the “diagonal effect” and “stimulated refill” liquidity process to a zero-intelligence model, thus gradually introducing strategic behaviour into the model.

Using nineteen days of data from the Paris Bourse in November 1991, [1] calculated the frequencies of events classified depending on aggressiveness¹, and our model builds on this classification. On the buy-side orders can be categorised into seven event types corresponding to differing levels of aggressiveness as detailed below (and summarised in Table 1).

A **Large Buy** event represents market orders that take liquidity through the best ask and through more than one price level by walking up the book. A **Market Buy** removes all orders at the best ask exactly. A **Small Buy** represents a market order to buy a quantity lower than that offered at the best ask. **New Bid Within** events are limit orders to buy posted within the best bid and ask quotes inside the spread. A **New Bid At** event posts a new bid limit order at best quotes joining the queue at the best bid and a **New Bid Away** posts a bid limit order at a price below the best bid. Finally a **Cancel Bid** refers to cancellations of bid limit orders irrespective of price level.

An additional seven event types can be defined analogously on the sell side of the book (storing the asks). Including off-book trades, we therefore have a total of fifteen different types of event. The results from the Paris Bourse study are summarised in Table 1 which shows the empirical probability distribution over the event types, with the probabilities adjusted to exclude applications². One important aspect to note is that their analysis is based on data that only gave the five best bid and ask price step levels of the book.

Bias et al. went on to analyse the same empirical probabilities over event types *conditional* on the previous event. The rescaled conditional probabilities from [1] are given in Table 2. Again we rescale their original results to exclude applications³.

¹By aggressiveness we mean the immediacy to which the agent is willing to trade.

²We do this as we are not interested in off-book trades. The probabilities are rescaled to sum up to 100%.

³The readjusting is again a simple rescaling to exclude applications so that the column probability vectors sum up to 100%. It does however ignore events at $t + 2$ if an application occurred at $t + 1$ e.g. the

Conditional frequencies of limit order book events														
(t-1)	Large Buy	Market Buy	Small Buy	New Bid Within	New Bid At	New Bid Away	Cancel Bid	Large Sell	Market Sell	Small Sell	New Ask Within	New Ask At	New Ask Away	Cancel Ask
Large Buy	7.51	3.51	15.80	13.45	5.16	6.08	5.37	1.85	1.14	14.83	7.48	4.79	7.04	5.99
Market Buy	3.26	2.52	17.45	2.18	7.02	11.55	8.08	3.42	6.28	23.95	2.83	2.50	4.94	4.02
Small Buy	3.80	2.96	21.00	9.91	5.40	6.49	4.12	2.41	1.79	20.26	6.03	4.65	6.35	4.84
New Bid Within	3.27	2.19	13.43	13.94	7.03	9.43	5.34	3.08	2.03	19.02	8.48	2.99	5.55	4.21
New Bid At	2.27	2.45	14.29	16.85	7.02	7.50	5.03	1.64	0.92	23.09	6.17	3.92	5.40	3.47
New Bid Away	1.94	1.63	11.72	7.65	8.14	18.59	4.54	3.29	2.46	20.38	7.05	3.44	5.69	3.49
Cancel Bid	2.33	1.76	10.85	12.79	5.35	8.27	9.97	3.90	2.32	20.08	8.94	3.67	5.95	3.82
Large Sell	0.94	0.95	6.62	9.05	4.66	6.70	6.12	10.09	3.42	24.52	12.89	3.42	5.12	5.51
Market Sell	2.93	6.41	16.37	2.31	2.96	4.34	3.65	4.88	3.08	29.73	2.24	5.04	9.62	6.45
Small Sell	1.62	1.63	10.49	7.07	4.88	5.93	4.47	4.93	3.02	34.61	8.05	4.09	5.41	3.79
New Ask Within	2.05	1.91	10.08	9.64	4.12	6.17	4.15	4.54	2.66	23.17	12.89	5.65	7.93	5.05
New Ask At	1.44	0.99	13.04	6.54	4.91	5.76	3.91	3.06	2.53	23.56	15.03	6.78	7.44	5.02
New Ask Away	2.29	2.2	11.04	7.39	4.21	5.98	3.43	2.87	2.14	21.80	5.97	7.22	19.77	3.68
Cancel Ask	2.08	1.90	10.20	8.95	4.50	6.25	4.62	4.16	2.18	19.67	14.19	4.05	7.76	9.47

Table 2: Conditional order and trade event frequencies rescaled from empirical results in [1, p. 1673]. Each entry in the main body represents a probability of an order book event (columns) conditioned by the previous event (rows). The three highest values in a given column are emphasised to illustrate the “diagonal effect”.

Book Side	Event Type	Probability(%)
Buy Side	Large Buy	2.4
	Market Buy	2.1
	Small Buy	12.7
	New Bid Within	9.2
	New Bid At	5.4
	New Bid Away	7.6
	Cancel Bid	4.8
Sell Side	Large Sell	3.9
	Market Sell	2.4
	Small Sell	24.5
	New Ask Within	8.5
	New Ask At	4.4
	New Ask Away	7.1
	Cancel Ask	4.9

Table 1: Absolute Frequencies of Limit Order Book Events rescaled from empirical results from the Paris Bourse [1, p. 1670].

As was noted in [1] the numbers on the diagonal tend to be larger than others in the same column. This “diagonal effect” suggests an event type is more likely to be preceded by the same event type than any other. Table 2 illustrates this along with the three most likely preceding events for each event highlighted in bold to show this diagonal effect.

3. STIMULATED REFILL

Along with the “diagonal effect” another interesting feature of order flow is the “stimulated refill” process ([2]) in order-driven markets where liquidity amongst liquidity providers and takers is self-organised to eliminate statistical arbitrage. In [1] they found most limit order activity being at or within quotes despite there being least depth at the best quotes when looking at the average profile of the order-book over time. They showed this to be a result of a cycle of transient depth whereby when spreads were tight, trades were more frequent thus widening the spread, followed by more within and at-quote limit order activity when spreads and depth at best bid or offer (BBO) were large. [1] and [2]

event chain: Large Market Buy (t) → Application (t+1) → Small Sell (t+2) is not rescaled as Large Market Buy (t) → Small Sell (t+1). Similarly it ignores any event sequence chains with a block of intermediary applications e.g. Event (t) → Application (t+1) ... → Application (T-1) → Event(T) etc. However we find this makes negligible differences to the probabilities.

argue that this behaviour is due to the liquidity replenishing and under-cutting/out-bidding activities of market-makers.

Table 3 below shows a stylised example set of probability vectors that show this effect of dynamic liquidity switching between regimes of submitting more trades and less orders versus a regime of submitting more orders and less trades dependant on the bid-ask spread. The probabilities are based on the empirical probabilities that [1] found when classifying the state of the book in a high or low spread regime in relation to the median value for the bid-ask spread.

In order to use these results in the context of a simulation model, several modifications are required. First all probability vectors are again rescaled to exclude applications which are non-order book events. Secondly when doing this particular analysis [1] grouped all their previous trade events classifications as just one trade event (e.g. all Large Buy, Market Buy, Small Buy events become classified just as a buy trade event). So in Table 3 we split the probabilities back into the more granular trade classifications by assuming they have the same relational proportions (amongst the three trade classifications on the relevant side of the book) as in the absolute case from Table 1. This maintains a probability for all fourteen events that will be needed for the model we develop in the next section. The last transformation averages their results that were based on further classifying the book as in a high or low depth regime as well as a high or low spread regime. This results in reducing their four probability vectors to two based simply on a high or low spread. As the fourteen order book event classifications from [1] only affect traded prices by shifting price levels of the order book, the volume depth at particular levels is not meaningful in the context of the model we will develop; therefore the issue of volumes of order book events is simplified in our analysis for now.

Thus in Tables 1, 2 and 3 we have three types of dynamics of order books that can be investigated by simulation of events based on the probability vectors.

4. THE MODEL

Our model is based on the event classification of the Bias et al. study [1] described in Section 2. Trading occurs over a series of discrete simulation steps. During each step an event type is chosen from Table 3, and we then simulate its effect on the order-book. We note any resulting trades and record the time series of transaction prices for later analysis. The

Event Type	Small Spread Probability (%)	Large Spread Probability (%)
Large Buy	3.0	2.0
Market Buy	2.6	1.7
Small Buy	15.7	10.3
New Bid Within	5.1	13.9
New Bid At	4.4	6.2
New Bid Away	7.0	8.5
Cancel Bid	4.9	5.0
Large Sell	4.5	2.9
Market Sell	2.8	1.8
Small Sell	28.5	18.1
New Ask Within	4.6	13.0
New Ask At	4.3	4.8
New Ask Away	7.2	7.4
Cancel Ask	4.7	4.5

Table 3: Dynamic frequencies of limit order book events (rescaled and modified from empirical results from [1, p. 1677] on the Paris Bourse in 1991).

event types are classified according to their liquidity-taking or liquidity-providing nature.

Our goal is to simulate this model under different experimental treatments, starting with a zero-intelligence treatment in which the type of event is chosen unconditionally at random from a discrete uniform distribution, followed by treatments which simulate more complex behaviour by allowing the probability with which a particular event is generated to be conditional on the current state of the market. This enables us to test the affect of conditional order flow adapting to either previous order flow or the state of the book; that is, to analyse the effect of *strategic* behaviour. Similarly to other models in the literature ([3, 9]), we use a volume size of 1 for all events except for the following: **Large Buy**, **Large Sell**, **Market Buy**, **Market Sell**. This is because these events are defined by removing a price step level of the order book at the best opposite quote and therefore need to have a volume corresponding to the available liquidity at best quotes. Prices of orders are rounded to a tick size of Δ^{Tick} and if the book is empty (e.g. at the start) and there is no best bid and offer (BBO) limit orders are posted with proximity to a reference fundamental price of P_f .

4.1 Market Events

In this section we detail how the simulated events introduced in Section 2 specifically affect the limit order book in our simulation model. We call b_t the best bid and a_t the best ask in the book at the current time of the simulation. All the event types can be classified as either market orders (**Large Buy**, **Large Sell**, **Market Buy**, **Market Sell**, **Small Buy**, **Small Sell**) or as limit orders (**New Bid Within**, **New Ask Within**, **New Bid At**, **New Ask At**, **New Bid Away**, **New Ask Away**, **Cancel Bid**, **Cancel Ask**)⁴. We now go through both of these two main order types detailing the model’s market mechanics for all event types when simulated in our model:

4.1.1 Market Order Events

Whenever a **Market Buy** event occurs, the volume at the best ask price is cleared out: all sell orders currently queued

⁴There are no marketable limit orders as limit order event types aren’t defined this way. In the original [1] analysis marketable limit orders would be counted as market orders.

on the order-book at this price step are removed (and vice-versa for **Market Sell** events at the best bid). This could still be a volume of 1 (similar to **Small Buy** and **Small Sell** events) if there is not a queue at the BBO. Trade events (**Large Buy**, **Market Buy**, **Small Buy**, **Large Sell**, **Market Sell** and **Small Sell**) occur if and only if there is available liquidity on the relevant book side. If not, the book remains unchanged and the simulation moves on to the next time tick.

For a **Large Buy**, a market order is submitted with volume equal to the total volume at the best ask +1 so that it removes liquidity at the best ask (a_t) and then trades at the next price level (a_t^2) above the best ask. A **Large Sell** event can be defined analogously taking liquidity on the bid-side of the book.

4.1.2 Limit Order Events

For limit orders submitted within the spread we assume they are posited randomly in between the spread i.e. uniformly deposited in the interval $[b_t, a_t]$. Limit orders posted at best quotes enter orders with a size of 1 at the end of the queue at BBO (b_t for **New Bid At** events, a_t for **New Ask At** events). Limit orders posted away from the spread are submitted uniformly anywhere between $[b_t^5, b_t]$ for **New Bid Away** events and $[a_t, a_t^5]$ for **New Ask Away** events where b_t^5 and a_t^5 represent the 5th best bid and 5th best ask level in the book. We do this as the probabilities are based on the empirical results of [1] who used order book data up to five price levels. If in the simulation there aren’t at least five price levels of liquidity in the book, the limit orders are posted uniformly in $[b_t-D, b_t]$ for bids and $[a_t, a_t+D]$ for asks where D represents a maximum offset parameter in the model. Limit order cancellations (**Cancel Bid**, **Cancel Ask**) remove at random an existing limit order in the book on the relevant side of the book. All limit orders on the relevant book side have equal probability of being deleted.

4.2 Experimental Treatments

We simulate our model under seven different experimental treatments. Each treatment is described in detail in the sections below and summarised in Table 4. The purpose of simulating the model under these different experimental treatments is to systematically evaluate whether incrementally introducing conditional behaviour into a zero-intelligence model produces more realistic time series data.

4.2.1 Equal treatment (“equal”)

Under this treatment, the event types are chosen randomly from a discrete uniform distribution over all fourteen event types. That is, the probability of generating any given type of event is $\frac{1}{14}$. This is the simplest treatment and corresponds to a pure “zero-intelligence” model.

4.2.2 Absolute treatment (“abs”)

Under this treatment the underlying distribution used to generate events remains static, but is calibrated with probabilities based on empirical data: that is, the probability of each event is given by the unconditional frequencies in the empirical analysis of Bias et al. on the Paris Bourse [1]. Thus the order book events are simulated with an unconditional discrete empirical probability distribution using the probability vector from Table 1. This corresponds to a calibrated “zero-intelligence” (ZI) treatment.

Section	Treatment Name	Experiment Conditions
4.2.1	“equal”	Equal Event Probabilities
4.2.2	“abs”	Absolute Event Probabilities
4.2.3	“cond”	Conditional Event Probabilities
4.2.4	“symmetric”	Symmetric Unconditional Event Probabilities
4.2.5	“dynamic”	Regime Probabilities based on fixed spread
4.2.6	“dynamicAlt”	Regime Probabilities based on localised window
4.2.7	“cluster”	Regime Probabilities alternating after random horizons

Table 4: Summary of experimental treatments

4.2.3 Conditional treatment (“cond”)

Under this treatment events are generated from a *conditional* discrete empirical probability distribution in which probabilities are conditioned on the previous event type. The probability vectors used are those in Table 2 and allows switching between 14 different probability distributions. For example if the last event simulated in the book was a **New Bid Within** the probability vector used to simulate the next event would be the 4th row in Table 2. The conditional order book event frequencies are rescaled from empirical results in [1, p. 1673]. This enables us to test the effects of adding the “diagonal effect” from [1]. Since behaviour in this model is conditional on previous order flow, we can ascribe to it a minimal level of strategic behaviour.

4.2.4 Symmetric treatment (“symmetric”)

Under this treatment events are generated from a static probability distribution, but with symmetric probabilities by averaging the unconditional empirical probability distribution to be symmetric on both buy and sell sides of the book. For example from Table 1 we can see in the “abs” case the probability for a **Large Buy** is 2.4% and for a **Large Sell** is 3.9%. In the “symmetric” case the probability of a **Large Buy** would equal the probability of a **Large Sell** with the value being $\frac{2.4\%+3.9\%}{2} = 3.15\%$. This is done with all 7 pairs of events that occur on the buy and sell side of the book.

4.2.5 Dynamic treatment (“dynamic”)

This treatment has probabilities conditioned on the state of the book choosing from the two spread regimes in Table 3. These dynamic frequencies are rescaled and modified from empirical results from [1, p. 1677]. It allows us to test adding the “stimulated refill” liquidity process to the unconditional ZI framework of the “abs” treatment whereby more liquidity taking occurs when spreads are low and more liquidity providing activity when spreads become high. We run this treatment after the “abs” treatment and use the average median spread value from that case as a fixed endogenous comparison value to switch between the high and low spread probability vectors depending if the spread at the current time tick is more or less than this comparison spread value respectively.

4.2.6 Alternative Dynamic treatment (“dynamicALT”)

This treatment is a variant of the “dynamic” treatment in which the comparison spread value alters and is calculated locally (from within the “dynamicALT” simulation run) as the median value from a localised rolling window of spreads in the last 100 simulation steps. Again, the probability dis-

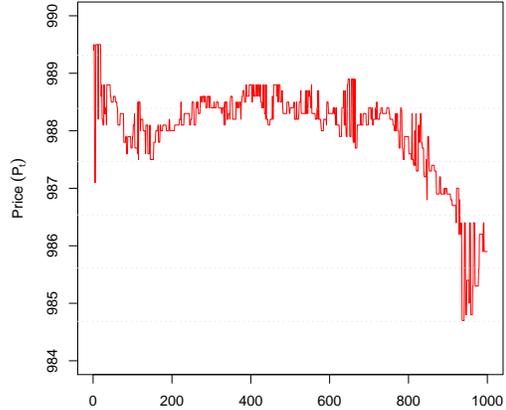


Figure 1: A sample price path for 1,000 time steps for the “abs” treatment

tributions are given by Table 3.

4.2.7 Cluster treatment (“cluster”)

In this treatment the dynamic switching between the two liquidity regimes is introduced exogenously in the simulation as the model simulates alternating between the two probability regimes (low to high to low etc..) from Table 3 with each regime lasting a L number of simulation steps where L is initialised at the start of each new regime by drawing a uniform random number between 1 and 100.

5. METHODOLOGY

For each of the above treatments we simulate the outcome on the order book of the events generated by the model.

At the start of the simulation there is no best bid or offer (BBO), therefore the initial market price at time $t = 0$ was set to the fundamental price $P_f = 10^3$. The tick size Δ^{Tick} was set to 10^{-1} and the maximum offset parameter $D = 5$ (which corresponds to 50 ticks). Each simulation was run for 1.1×10^4 time steps. We recorded prices only from $t > 10^3$ onwards in order to wash out any initial effects.

We used the Mersenne Twister ([10]) algorithm to draw all random variates in the simulation. We ran a total of seven experimental treatments which are summarised in Table 4. For each treatment we ran 10^4 independent simulations, thus producing a total of 7×10^4 different time series of prices and returns, which were then analysed for the stylized facts.

Figure 1 shows a typical section of the time series of prices

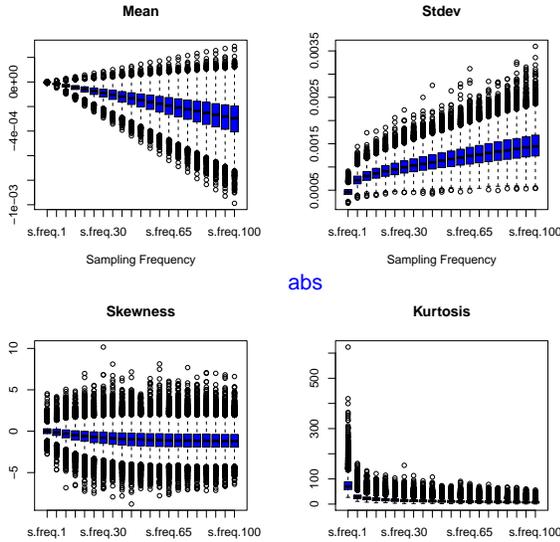


Figure 2: Box-Whisker plot for the Mean, Standard Deviation, Skewness and Kurtosis of returns from the “abs” treatment.

produced over the course of 10^3 simulation steps. From the price time series we calculated the time series of returns r_t according to the following equation:

$$r_{t_1} = \log(p_{t_1}) - \log(p_{t_0}), \quad (1)$$

where p_t is the market price (defined as the last trade price) at time t as measured in discrete “event time”: that is, the time value is incremented by 1 each time an order book event occurs and represents the number of simulation steps that have elapsed since the start of the simulation. Prices are sampled at regular time intervals $t_1 - t_0$.

We then tested for the presence of the following statistical features in the simulated time series of returns as given by equation 1: fat tails, volatility clustering, persistence and non-Gaussianity. For each treatment we computed the (a) Hill-estimator, (b) Hurst exponent, (c) p -values for the ARCH-LM test and (d) p -values for the Jarque-Bera test. For the Hill-estimator the 0.025, 0.05, 0.95 and 0.975 quantiles were chosen. Because the return calculation depends on the frequency with which prices are sampled ($t_1 - t_0$), the tests were calculated for each sampling frequency in the set $\{5, 10, \dots, 95\}$ as measured in simulation steps. A distribution of stylised fact test results is obtained from the 10^4 simulation runs which allows us to run t -tests comparing treatments.

6. RESULTS

We first show some basic statistical results from the returns on the “abs” treatment in Figure 2. We can see the mean returns are negative, as one might expect given the asymmetry in the empirical probability distribution: as can be seen from Table 1 sell side of the book events have higher probabilities than the corresponding buy side of the book events reflecting the empirical data collected in [1]. Again, as one might expect, as the sampling interval increases the re-

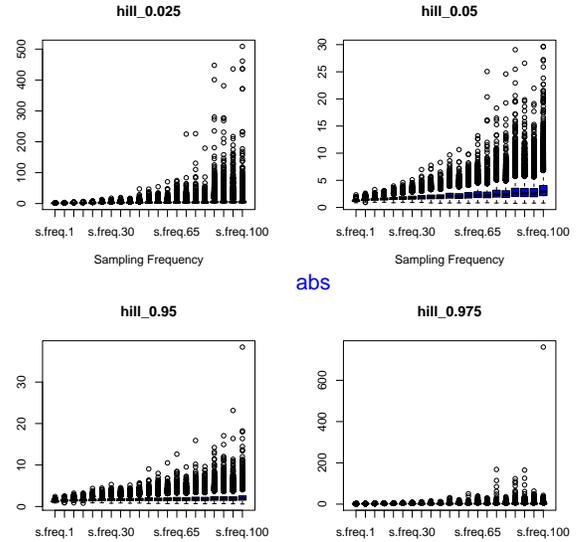


Figure 3: Box-Whisker plot for the Hill Estimator of returns from the “abs” treatment at the lower (0.025, 0.05) and upper (0.95, 0.975) quantiles.

turns exhibit higher standard deviations. We observe slight negative skewness and significant kurtosis suggesting that the “abs” treatment reproduces non-Gaussianity and fat tails.

This is confirmed when looking at the Hill estimator in Figure 3 which tends to around 3 in the upper and lower tail. In Figure 4 we can see we get a Hurst exponent ≈ 0.6 which [9] and [12] argue to be realistic on short to medium timescales. From Figure 4 we can see the ARCH-LM test shows no evidence of volatility clustering (p -values much higher than a 5% significance level) but the Jarque-Bera test shows significant deviance from Gaussian returns (p -value close to 0).

We now turn to the results of performing the same statistical tests on the other treatments. Figure 5 compares the sample means of summary statistics run on the returns of the seven treatments described in Section 4.2 and summarised in Table 4. We can see in Figure 5 that apart from the “equal” and “symmetric” treatments all other treatments have a negative return due to asymmetry with higher probabilities for events on the sell side of the order book.

With the exception of the “equal” treatment all the treatments show significant kurtosis and reproduce non-Gaussianity in returns. Figure 6 shows the sample means of the Hill estimator tests on the 7 treatments and again shows with the exception of the “equal treatments” all other treatments show similar behaviour reproducing fat tails in the returns. Figure 7 compares the Hurst Exponent, ARCH-LM test p -values and Jarque-Bera test p -values of returns for the 7 treatments and show all models reproduce realistic high frequency Hurst exponents ranging from 0.4 to 0.6. The ARCH-LM test shows like the “abs” treatment none of the treatments reproduce volatility clustering in returns. The Jarque-Bera test p -values confirm that with the exception of the “equal” treatment the other treatments exhibit non-Gaussian returns.

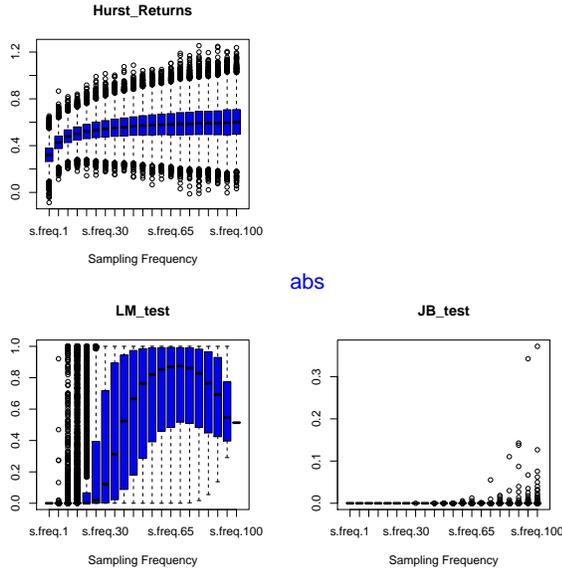


Figure 4: Box-Whisker plot for the Hurst Exponent, ARCH-LM test p -values and Jacques-Bera test p -values of returns from the “abs” treatment.

7. DISCUSSION

In all experiments other than the “equal” treatment, we observe realistic Hurst exponents, fat tails and non-Gaussianity in return behaviour. Thus our results show that *unconditional* zero-intelligence order flow is sufficient to reproduce these phenomena, provided that the probability distribution used to drive the zero-intelligence behaviour is calibrated against empirical frequencies (as per the “abs” treatment).

In the treatments “cond”, “dynamic”, “dynamicALT” and “cluster” we have simulated *conditional* order flow (a proxy for strategic behaviour) and found similar results to the unconditional order flow treatments. None of our experiments reproduced volatility clustering. This suggests that a more complex correlated and organisation in order flow occurs in real markets.

8. CONCLUSION AND FUTURE WORK

In this paper we developed a zero-intelligence model of financial markets based on an empirical analysis of the Paris Bourse exchange [1]. We used this model to test whether a “zero-intelligence plus” model could simultaneously explain several of the stylized facts of financial time series data by systematically evaluating seven different experimental treatments in which we incrementally added minimalist levels of strategic (conditional) behaviour to our model. We tested the resulting time series of returns for the following statistical features: fat tails, volatility clustering, high-frequency persistence and non-Gaussianity. Surprisingly, the “zero-intelligence plus” treatments do *not* introduce more realism to the time series.

Our new zero-intelligence model therefore highlights the fact there are many ways to simulate random order flow and within each structure there is implicit intelligence. In our model agents specifically remove or add liquidity to the market as the random order book events are classified with implicit knowledge about price levels in the order book. In the [9], [11], [8] ZI models, it is left to the matching of the double auction to determine whether liquidity is removed or

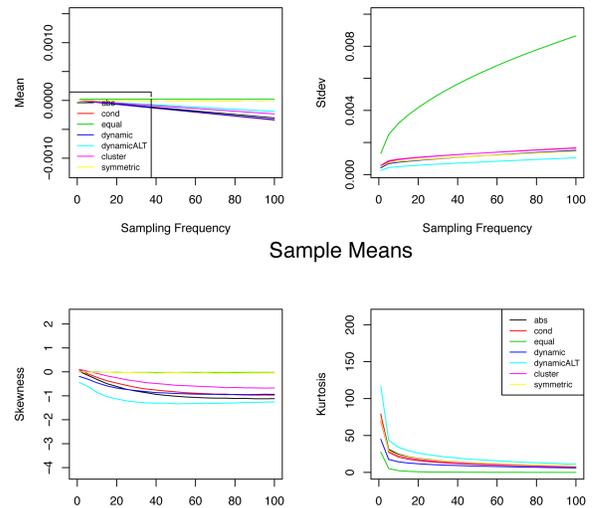


Figure 5: Sample means for mean, standard deviation, skewness and kurtosis of returns for the seven treatments.

added changing price levels in the book despite the fact in a continuous double auction market with a visible order book this would be known to all agents. This motivates setting up a simulation model which takes into account more knowledge of the state of the book known in real time.

Using the decomposed microstructural events of [1] in our model, a new framework can be used to investigate adaptive order flow. Specifically it allows a more detailed understanding of the level of organised order flow amongst liquidity providers and takers is needed to understand real markets. We find changing the model parameters alters the degree of reality and introducing calibrated vectors improves results of stylised facts in terms of realism. Similar to other ZI models in the literature, the model manages to reproduce some realistic stylised facts, thus supporting the claim they are due to the microstructure and not strategic behaviour. In this case we manage to reproduce realistic Hurst exponents, fat tails and non-Gaussianity of returns. The model however in all treatments did not produce volatility clustering suggesting there is a more complex and correlated interplay of orderflow required for this phenomena. In the “cluster” treatment of our model a simple clustering of transactions was not sufficient to reproduce volatility clustering. Our results thus support the claim by [6] and [2] that along with ordering of transactions, there are also subtle fluctuations in the balance between liquidity taking and provision that acts as important factors in determining clustered volatility. This they argue would need to incorporate long-memory of supply and demand incorporating both order splitting and “stimulated refresh” liquidity provision. [14] discuss non-ZI explanations of volatility clustering suggesting social interactions amongst agents which employ herding behaviour (either by direct or indirect imitation) and/or order splitting of large trades are related to this phenomena.

The model could be extended with more complex dynamics, for example by using a reinforcement learning algorithm in which the propensity to choose a particular type of event adapts over time according to a profit signal such as the ef-

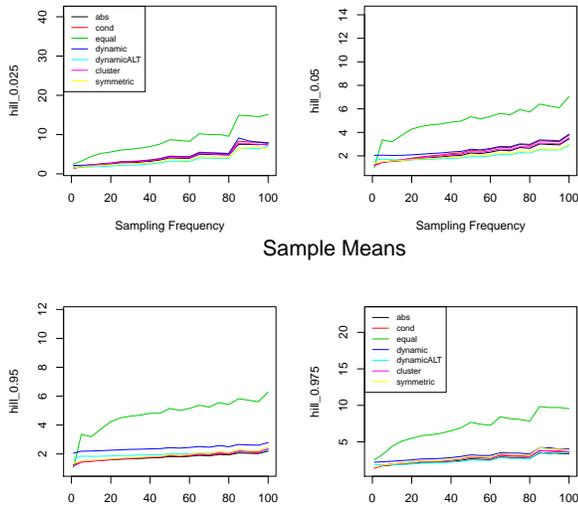


Figure 6: Sample means for the Hill Estimator of returns for all 7 treatments at the lower (0.025, 0.05) and upper (0.95, 0.975) quantiles.

fective cost of an order type. Such a learning model could possibly induce a more realistic “stimulated refresh” effect whereby clustered buying (selling) market order activity is followed by clustered selling (buying) limit order liquidity replenishing. Investigating this effect could provide an interesting avenue for future research to see its effect in reproducing realistic heavy tail and volatility clustering phenomena.

Our results provide further evidence that some phenomena in financial markets can be attributed solely to the market design, and do not require strategic behaviour. However, our findings also highlight the importance of time correlation of order flow; markets exhibit a complex interplay of latent supply and demand arising from strategic behaviour in which agents adapt their liquidity-taking and liquidity-providing behaviour depending on other agents’ actions.

9. REFERENCES

- [1] B. Biais, P. Hillion, and C. Spatt. An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse. *The Journal of Finance*, 50(5):1655–1689, 1995.
- [2] J. P. Bouchaud. The Endogenous Dynamics of Markets: Price Impact, Feedback Loops and Instabilities. In A. Berd, editor, *Lessons from the Credit Crisis*. Risk Publications, 2011.
- [3] C. Chiarella and G. Iori. A simulation analysis of the microstructure of double auction markets. *Quantitative Finance*, 2(5):346–353, 2002.
- [4] D. Cliff and J. Bruten. Minimal-Intelligence Agents for Bargaining Behaviors in Market-Based Environments. Technical report HPL-07-91, HP Labs, 1997.
- [5] R. Cont. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2):223–236, Feb. 2001.

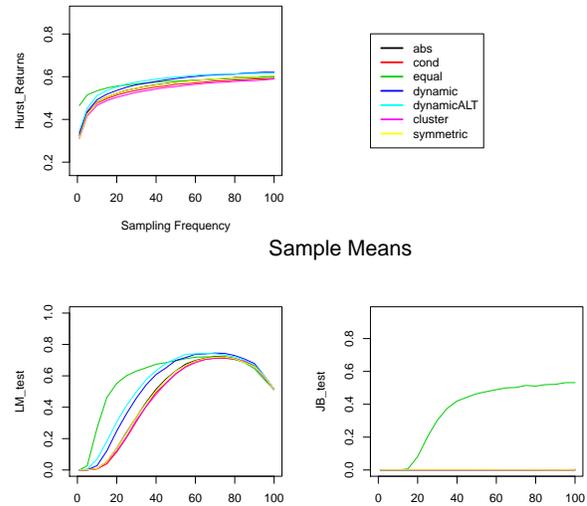


Figure 7: Sample means for the Hurst Exponent, ARCH-LM and Jacques-Bera test p -values of returns for all seven treatments.

- [6] L. Gillemot, J. D. Farmer, and F. Lillo. There’s more to volatility than volume. *Quantitative Finance*, 6(5):371–384, 2006.
- [7] D. K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137, 1993.
- [8] D. Ladley and K. R. Schenk-Hoppé. Do stylised facts of order book markets need strategic behaviour? *Journal of Economic Dynamics and Control*, 33(4):817–831, Apr. 2009.
- [9] S. Maslov. Simple model of a limit order-driven market. *Physica A: Statistical Mechanics and its Applications*, 278(3-4):571–578, Apr. 2000.
- [10] M. Matsumoto and T. Nishimura. Mersenne twister: a 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Transactions on Modeling and Computer Simulation*, 8(1):3–30, Jan. 1998.
- [11] S. Mike and J. D. Farmer. An empirical behavioral model of liquidity and volatility. *Journal of Economic Dynamics and Control*, 32:200–234, 2008.
- [12] T. Preis, S. Golke, W. Paul, and J. J. Schneider. Multi-agent-based Order Book Model of financial markets. *EPL (Europhysics Letters)*, 75(3):510–516, Aug. 2006.
- [13] M. Wyart, J.-P. Bouchaud, J. Kockelkoren, M. Potters, and M. Vettorazzo. Relation between bid-ask spread, impact and volatility in order-driven markets. *Quantitative Finance*, 8(1):41–57, 2008.
- [14] R. Yamamoto. Volatility clustering and herding agents: does it matter what they observe? *Journal of Economic Interaction and Coordination*, 6(1):41–59, May 2011.