

MAGKS



**Joint Discussion Paper
Series in Economics**

by the Universities of
**Aachen · Gießen · Göttingen
Kassel · Marburg · Siegen**

ISSN 1867-3678

No. 43-2014

Alexandru Mandes

**Order Placement in a Continuous Double Auction Agent
Based Model**

This paper can be downloaded from
http://www.uni-marburg.de/fb02/makro/forschung/magkspapers/index_html%28magks%29

Coordination: Bernd Hayo • Philipps-University Marburg
Faculty of Business Administration and Economics • Universitätsstraße 24, D-35032 Marburg
Tel: +49-6421-2823091, Fax: +49-6421-2823088, e-mail: hayo@wiwi.uni-marburg.de

Order Placement in a Continuous Double Auction Agent Based Model

Alexandru Mandes*

University of Giessen

Licher Str. 64

35394 Giessen, Germany

email: Alexandru.Mandes@wirtschaft.uni-giessen.de

August 20, 2014

Abstract. Modeling intraday financial markets by means of agent based models requires an additional building block which reflects the order execution, i.e. the trading process. Current implementations rely only on stochastic placement strategies, ranging from total randomness to adding some budget constraints. This contribution addresses the issue of order placement for low-tech traders, by replacing the zero-intelligence assumption with a microtrading-based approach. The results show that the power-law decaying relative price distribution of off-spread limit orders and the concave shape of the overall market price impact can be replicated when rational order submission strategies are used.

Key Words: agent based modeling, high-frequency financial markets, continuous double auction, order placement, market impact

JEL classification: C63, N20

*The author would like to thank Deutscher Akademischer Austauschdienst (DAAD) for the awarded PhD scholarship.

1 Introduction

A literature review identifies two main classes of Agent Based Models (ABMs) for financial markets with respect to their time-line frequency. On one side, *daily ABMs* focus on the behavior of individual agents, but simplify to a great extent the process of price discovery. The daily market equilibrium price is updated usually either through a market impact function, where the excess demand is cleared by the market maker at the new adjusted price (e.g. Lux (1995, 1998), Lux and Marchesi (1999, 2000), Chen and Yeh (2001, 2002), Farmer and Joshi (2002), Westerhoff (2009, 2010)), or by means of a Walrasian tâtonnement mechanism of identifying the equilibrium where there is no aggregate excess demand (e.g. Arthur, Holland, LeBaron, Palmer and Tayler (1996), Brock and Hommes (1997), Fischer and Riedler (Forthcoming)). Both approaches are highly synchronous, i.e. agents' demands are batched together and executed at the same time and at the same price. These low-frequency ABMs describe the emergent market dynamics by a single price, sometimes associated with a total volume.

A different class of ABMs zooms into the intraday world where trading takes place within a continuous double auction framework. These models are also referred to as *disequilibrium models* since agents asynchronously disclose their trading demands, and transactions can take place outside the equilibrium price. Price changes can be seen as an outcome of the interplay between order flow and the persistent order book liquidity. Moreover, their dynamics is characterized by tick-by-tick data on order flow (issuing new orders, modifying, deleting, expiration of limit orders), quotes (bid-ask spread and order book shape) and trades (price and volume).

In contrast to daily ABMs, an additional decision layer dealing with order execution is required. Most intraday ABMs implement only stochastic placement strategies, ranging from pure randomness such as in Cui and Brabazon (2012) to adding budget constraints such as in Chiarella, Iori and Perelló (2009).¹ The *zero-intelligence* approach represents the most appealing way of circumventing a difficult problem and, from an historical perspective, is one of the earliest solutions proposed.² Moreover, a model with random agents allows for the assessment of market institutional design such as in Farmer, Patelli and Zovko (2005) and can also provide a benchmark for other ABMs involving more rational agents. On the other side, besides their lack of realism – the “promise” of ABM is to provide a sound micro-based design –, these models are confronted with several limitations.

¹The price is randomly drawn from an interval restricted by the current budget and return expectation. This stochastic limit price and the current portfolio structure determine the sign of the individual order as well as its size.

²Similarly, in the case of the traditional daily ABMs one of the first order formation (investment) strategies proposed by Gode and Sunder (1993) was based also on random behaviour. The authors are also the ones who have actually coined the “zero-intelligence” term.

For example, Cui and Brabazon (2012) conclude that replicating a realistic price impact of market orders cannot be achieved without agent intelligence. In real markets, trade size and timing are not random, but rather take into account the existing market liquidity – just by inspecting the depth of the order book a large market impact can usually be avoided when execution time permits. If this liquidity factor is ignored, the market impact in a simulation experiment is higher for larger orders than in the case of real markets, such as replicated in Cui and Brabazon (2012).

We also stress that the order flow generated by these low-frequency agents, as well as the eventual shape of the order-book, play key roles in building more complex intraday market models, where further microstructure-based trading strategies – such as algorithmic traders, market makers or other high-frequency traders – rely on these sources of information. Thus, simplifying too much the way low-frequency agents execute their orders influences the general intraday environment and could further on affect the behavior of high-frequency traders.

This contribution addresses the issue of order placement for low-tech strategies by replacing random trading decisions with a liquidity and volatility-based optimisation approach. We introduce more intelligence in order execution, by taking into account the current market state as well as intrinsic agent characteristics, and inspect how the intraday market dynamics changes.

Section 2 starts with the description of the order placement problem in subsection 2.1. We briefly mention three modelling approaches present in the literature and also introduce how our model relates to them. Following, a set of microstructure factors and their relationship to various order properties are described. In subsection 2.2 we present in more detail our model’s components, parameters and assumptions. An iterative numerical procedure for identifying the optimal relative limit distance is presented in subsection 2.3. Section 3 describes how the order submission model is integrated within an agent based model framework. In section 4, we present the experimental results benchmarked against a zero-intelligence model with respect to the relative price distribution of limit orders and price impact of market orders. Finally, we present our conclusions and further research options.

2 Order placement in a continuous double auction

The *double auction* is one of the most common mechanisms of price discovery in equity markets. Basically, participants can place their trading offer and demand as *market* or *limit orders*. New orders are matched against an




Figure 1: Order book illustration

outstanding order book formed of two queues of limit orders – one for buy (*bid*) and one for sell orders (*ask*) – which are ordered by price and time priority as in figure 1. Prices are discretised up to a minimum increment called *tick*. The highest bid price (296.98) and the lowest ask price (297.30) are called *best quotes* and the difference between them (0.32) is known as *bid-ask spread*. Moreover, the centre of the spread is referred to as the *midpoint* or *midprice* (297.14). An incoming market order is sequentially executed against the available limit orders on the other side of the market, ordered by their priority, until the entire order size is filled. On the opposite, a new limit order which does not cross any outstanding limit orders is stored in the order book at the specified price and waits to be executed against future arriving market orders. A graphical representation where some related concepts are identified is provided in figure 1. Two key measures which will be extensively used in the rest of this paper are the *relative limit distance* Δ , i.e. the difference between the limit price and the best quoted price on the opposite side of the market, and the *relative limit price* δ , defined in Zovko and Farmer (2002) as the difference from the best quote on the same side. The two measures differ by an amount equal to the current spread.

2.1 Order placement problem

The trading process is a trade-off between execution cost and delay risk and comprises two types of trading decisions: *order scheduling* (break-up large orders or all-in-one piece) and *order submission* (type choice and placement). In this paper we will tackle only the later one, which is actually the foundation for designing a trading-driven ABM. Proper execution of individual orders involves a set of micro-trading decisions, i.e. the order can be articulated into a market order, a limit order or can be split between a market order and a limit order. In other words, the trader faces a trade-off between execution certainty and a more favourable transaction price. At one extreme, a

market order does not carry any execution risk, but has a higher transaction cost consisting in market impact. On the other side, preferring a limit order saves the cost of immediacy associated with the market order alternative and can further improve trading costs with the relative limit distance. The drawback is that limit orders encounter the risk of remaining unexecuted, as they are conditioned on the uncertain future event of being matched by a counter party.

Several papers have studied various order placement strategies for limit order book markets. Lillo (2007) defines an optimisation problem within the framework of expected utility maximisation, where the probability for a limit order of getting executed is given by the first passage time distribution of a random walk, i.e. the probability that the stochastic price reaches the limit price Δ by a certain time – thus, the hitting time probability is a function of Δ , time horizon and volatility. Kovaleva and Iori (2012) develop a model which discriminates between placing a market order and a limit order at an optimal limit price. In their model, the total time-to-fill is not only given by the first passage time distribution which measures the probability of reaching the beginning of the queue, but also by a random delay – sampled from an exponential distribution with constant intensity – which stands for the order’s effective execution. By assuming stochastic log-normal processes for the trajectories of the bid and ask prices, an analytical solution is identified by maximizing a mean-variance utility function.

Cont and Kukanov (2012) formulate a convex optimisation problem for the decision of splitting between a market and a limit order placed at the best bid or ask. The optimisation function penalises the execution price with an execution risk which takes into consideration factors such as order size, the existing queue size at the front of the book and the cumulative distribution function of the queue outflow. If a functional form for the outflow CDF is assumed, a parametric numerical solution can be computed or, alternatively, a non-parametric solution when the empirical distribution is based on past order fills. The authors also extend the allocation problem into a routing problem across multiple trading venues, further taking into consideration liquidity fees and rebates structure.

Our contributed model is more related to Cont and Kukanov (2012) by the construction of the optimisation problem – deciding with respect to the market-limit order split with the goal of minimising the risk adjusted execution cost. One distinction is that limit order placement is not restricted any more to only the best quotes and thus a supplementary decision about the optimal limit price has to be made. More important is that our model relies only on intrinsic agent characteristics and several microstructure factors. The latter can be directly observed and assessed within an agent based framework with a limit order book mechanism, without assuming any functional forms of the underlying price or order flow processes. The agents’

order placement strategies are not optimal in the sense of being derived under expectations, but appear as rule-of-thumb strategies expressing a form of bounded rationality. The strategy inputs reflect different facets of the current market state making agents more reactive, in the spirit of agent-based modeling where reflexivity is a key concept.

Johnson (2010) provides a broad overview of factors driving traders to act more aggressively (impatient) or passively (patient), which can be classified into liquidity-, price- and time-based factors. Firstly, the choice in favour of a market order is found to be highly dependent on the instant liquidity reflected by the market tightness, i.e. the bid-ask spread, and by the order book height, i.e. the potential price impact of a market order walking up the book. Aggressive orders are more probable when the cost of immediacy is low, while a high liquidity cost due to higher spreads encourages liquidity suppliers to place more limit orders. Beber and Caglio (2005) found also evidence for a non-linear relation showing that particularly wide spreads favour in-spread limit orders rather than market or far away off-spread orders. Pascual and Veredas (2009) concluded that wide spreads discourage especially small market orders, increasing the frequency of larger market orders. With respect to order size in general, Cont and Kukanov (2012) state that market orders are usually larger than limit orders.

Order book depth influences the general order aggressiveness in two ways. An overall supply-demand imbalance drives traders to price their orders more aggressively when their side of the book is thicker and crowded in order to increase their order execution probability (competition effect). Conversely, traders become less aggressive when the opposite side is deeper, forecasting a favorable short-term order flow (strategic effect). If only the thickness of the opposite side of the market is taken into consideration, both Beber and Caglio (2005) and Pascual and Veredas (2009) identify an asymmetric behavior – sellers are more impatient to trade than buyers and thus are more willing to take advantage of the available liquidity by issuing larger aggressive orders; contrary, buyers show more patience and place less aggressive orders.³

Zovko and Farmer (2002) found that short-term volatility at least partially drives the relative limit prices and also suggest that such a feedback loop may contribute to volatility clustering. On one side, the probability of execution for further placed limit order increases and, on the other side, the picking-off risk due to adverse selection is also higher. Another price-based factor is the momentum indicator proposed by Beber and Caglio (2005), defined as the ratio between the current price and its exponential moving average. The direction of the short-term market trend asymmetrically affects the execution probability of limit orders and ultimately leads to a change in order pricing aggressiveness. Moreover, higher previous traded volumes, acting as a proxy

³Sellers are more concerned about the non-execution risk, while buyers pay more attention to the picking-off risk due to misspricing their orders.

for market information, lead to an even bigger increase in aggressiveness in the direction of the market trend.

2.2 Order placement model

The model proposed in this paper is set up as an *optimisation problem* which consists in making two decisions: (i) discriminating between a market and a limit order, and (ii) – in the case of a limit order – identifying the optimal relative limit distance which minimises the risk adjusted execution cost. Several liquidity- and price-based components are included.

The *objective function* $f(M, \Delta)$ describes the trade-off between execution cost and non-execution risk, balanced by agent's sense of urgency λ_u . The *sense of urgency* can reflect a mix of risk aversion, degree of informativeness, strategy time-frame or just time pressure, i.e. waiting time of the trading process. The decision variables are the *fraction* M of the total order executed as a market order and the *relative limit distance* Δ for the outstanding quantity $(1 - M)v$ traded as a limit order.⁴

$$\min_{M, \Delta} f(M, \Delta) \quad (1)$$

$$f(M, \Delta) = \text{cost}(M, \Delta) + \underbrace{\lambda_u \text{risk}(M, \Delta)}_{\text{adjusted risk}} \quad (2)$$




Figure 2: Examples of buy order aggressiveness, relative to the base price

The *cost function* $\text{cost}(M, \Delta)$ captures what is known as the *implementation shortfall* $\text{imp.sh}(M, \Delta)$, i.e. the difference between a given benchmark BM_{is}

⁴The function is not identified for $M = 1$, reflecting the situation in which the entire quantity is traded as a market order, which is logical consistent with the fact that a decision regarding Δ is no longer required.

and the effective order execution price. One of the most common benchmarks in liquid markets is the arrival price, i.e. the current bid-ask midpoint, but also other benchmarks, such as the last trade price or previous day close, can be considered.⁵ For example, in the extreme case when the entire quantity is executed as a market order ($M = 1$), the implementation shortfall reflects exactly the price of immediacy – equal to half the spread – plus any additional market impact. The *market impact function* $\text{mk.imp}(v)$ is influenced by the current order book state and can be computed as the percentage change in price where the entire order size is executed. Actually, all measures involved (Δ , BM_{is} , mk.imp) are scaled as percentage returns relative to a *base price*,⁶ which is the best bid for sell orders or the best ask for buy orders, respectively.⁷ An example of how these measures are related for different types of buy orders is depicted in figure 2. The order volume V is usually expressed as a percentage of the average daily volume (ADV).

$$\text{imp.sh}(M, \Delta) = \underbrace{M (BM_{is} + \text{mk.imp}(MV))}_{\text{market order part}} + \underbrace{(1 - M) (BM_{is} - \Delta)}_{\text{limit order part}} \quad (3)$$

$$= BM_{is} + M \text{mk.imp}(MV) - (1 - M) \Delta \quad (4)$$

The cost function in equation (5) wraps around the implementation shortfall by adjusting it with a *volatility-threshold downside price change penalty* in order to discourage execution prices which are too far away beyond a multiple σ_{is} of the short-term volatility, corresponding to highly unfavourable executions.⁸ Besides exponent θ , the parameter $\beta > 1$ controls for the size of this penalty.

$$\text{cost}(M, \Delta) = \begin{cases} \text{imp.sh}(M, \Delta) & \text{for } \text{imp.sh}(\cdot) \leq \sigma_{is} \\ \beta \sigma_{is} (\text{imp.sh}(M, \Delta) / \sigma_{is})^\theta & \text{for } \text{imp.sh}(\cdot) > \sigma_{is} \end{cases} \quad (5)$$

On the risk side, the *execution probability* of a given limit order depends on

⁵In the arrival price case, the relative execution benchmark $BM_{is} = s/2$, where s is the bid-ask spread.

⁶Because of the additivity property, log returns would have been more precise in computing the differences between the aforementioned measures, but the derivations in subsection 2.3 would have become analytically intractable. Eventually, as we are dealing with small intraday deviations, the imprecisions associated with the use of simple percentage returns appear acceptable.

⁷In case the bid and/or ask values are missing, the last trade price can substitute as base reference, as well as execution benchmark.

⁸If the return expectation of the agent is known, an execution threshold taking into consideration also this value could be implemented.

(i) *short-term market volatility* $\text{dyn}(\Delta)$,⁹ (ii) *order flow* proxied by the order book imbalance $\text{flow}(OBI)$,¹⁰ and (iii) *order queue* in front of the limit order queue $\text{queue}(\Delta)$. Moreover, (iv) an *opportunity cost* as a penalty function of order size $\text{size}(\nu)$ is also included. The aggregate non-execution risk function $\text{risk}(M, \Delta)$ is given by (6). The parameters $\alpha_0, \alpha_1, \alpha_2$ tune the general preference for market and limit orders, as well as the distribution of relative limit distances. Eventually, each individual decision will also depend on the agent’s urgency factor λ_u and on the specific state of the market.

$$\text{risk}(M, \Delta) = (1-M) \text{flow}(OBI) \text{size}((1-M)V) (\alpha_0 + \alpha_1 \text{dyn}(\Delta) + \alpha_2 \text{queue}(\Delta)) \quad (6)$$




Figure 3: (left) Market dynamics effect as a function of Δ and ω with fixed $BM_{dyn} = 1\%$, $\sigma_{dyn} = 1.5\%$; (right) Order-flow as a function of μ and OBI

The functional form of the *market dynamics effect* $\text{dyn}(\Delta)$ describes a sub-linear increasing risk of non-execution for limit orders inside the volatility bands, defined by a central benchmark BM_{dyn} and a multiple of the short-term volatility σ_{dyn} . When the relative limit distance is outside this interval, the risk increases faster penalising far away orders.¹¹ In the same spirit, a larger exponent ω increases the relative weight associated with the limit distance effect. Potential candidates for BM_{dyn} can be for example the bid-

⁹According to Johnson (2010), short-term or transient volatility is mostly liquidity-driven, while fundamental volatility is more long-term and caused by informational shocks.

¹⁰Order flow can also be computed as the imbalance in the value or number of incoming orders over a period of time.

¹¹The intuition behind this effect could be seen as similar to the technical trading tool known as “Bollinger Bands” (see Bollinger (2001)), which relies on the price dynamics fluctuating inside an interval bounded, under standard parameters, by two standard deviations above and below a 20 periods (days) moving average.

ask midpoint, last trade price or previous close price.

$$\text{dyn}(\Delta) = \sigma_{\text{dyn}} \left(\frac{\Delta - BM_{\text{dyn}}}{\sigma_{\text{dyn}}} \right)^\omega \quad (7)$$

A key model component is the *expected order-flow flow(OBI)*, which drives the short-term price returns and affects the execution probability of outstanding limit orders. In this implementation, order flow is more liquidity- rather than price-based and relies on the order book imbalance (OBI) indicator. *OBI* quantifies the difference between the cumulated volumes up to a certain depth level N on each side of the order book.¹² By its definition in (9), *OBI* takes values between -1 and 1 , the extremes corresponding to the cases when one book-side is empty. When the order book is unfavorable leaned, *OBI* is positive, $\text{flow}(OBI)$ is larger and the ultimate non-execution risk increases leading to a bigger incentive of placing more aggressive limit orders. Overall, $\text{flow}(OBI)$ takes values between $1/\mu$ and μ – asymmetric around zero – meaning that it also acts as a complementary penalty factor when *OBI* is unfavorable. The constant μ in equation (8) weights the relative importance of the order flow effect in assessing the aggregate non-execution risk and in setting the optimal relative limit distance Δ , e.g. *OBI* is neutral when $\mu = 1$. Finally, it is useful from the implementation perspective if *OBI* has a different sign for buy and sell orders.

$$\text{flow}(OBI) = \mu^{OBI} \quad (8)$$

$$OBI = (-1)^{\mathbb{1}_{\text{sell}}} \frac{\sum_{i=1}^N \text{bid}_i - \sum_{i=1}^N \text{ask}_i}{\max(\sum_{i=1}^N \text{bid}_i, \sum_{i=1}^N \text{ask}_i)} \quad (9)$$

The effective *order queue effect queue*(Δ) reflects the cumulative size of the book queue BQ_Δ situated in front of the client limit order – placed at the relative distance Δ . The order queue can be immediately computed within an observable order book and is expressed directly as a percentage of ADV, without assuming any functional form.

$$\text{queue}(\Delta) = BQ_\Delta \quad (10)$$

¹²Alternatively, the book imbalance and the short-term price return can be assessed by comparing the bid-ask midpoint with the weighted price at a given depth of the order book as in Cao, Hansch and Wang (2009).

The *opportunity component* $\text{size}(\nu)$ is an increasing function of order size, always bigger than one because of the exponential. Since order size ν takes most of the time subunitary values very close to zero, the exponent η should also be less than one in order to be able to discriminate between the various order sizes. The intuition behind this penalty is that an outstanding limit order is associated with a “*signaling*” risk as well as a “*jump-over*” effect – the bigger the order, the more likely other limit orders get placed in front. Furthermore, a non-executed limit order is expected to be transformed into a market order at a worse transaction price than the initial one, because the market is assumed to have moved in an unfavorable direction, i.e. *adverse selection*.

$$\text{size}(\nu) = \exp(\nu^\eta), \tag{11}$$

where, in the general case, $\nu = (1 - M) V$.

2.3 Order placement strategy

There are three issues in trying to analytically deal with the optimisation problem defined in subsection 2.2: (i) in the case of continuous $0 \leq M \leq 1$ the resulting exponential equation due to (11) can only be solved by applying the Lambert W function (omega function) or numerical procedures; (ii) unless a functional form for the order queue effect $\text{queue}(\Delta)$ is assumed, an analytical solution cannot be derived; (iii) certain polynomial degree restrictions with respect to exponents θ and ω need to be made in order to have unique solutions.

We are willing to partially simplify the minimisation problem in equation (1) by restricting M to take only binary values 0/1, which corresponds to choosing only between a market and a limit order for the entire quantity.¹³ Furthermore, we restrict $\theta = \omega = 2$, as there is no need for a higher degree penalty – any necessary tuning is possible by adjusting the α_i and β parameters. However, if a functional form for $\text{queue}(\Delta)$ based, for example, on an average order book shape would be assumed, any connection to the temporal specific structure of the book would be lost. Therefore, we implement an *iterative numerical procedure* for identifying the optimal relative distance Δ^* of a potential limit order. The implications of these decisions are presented in the rest of this subsection.

As a consequence of restricting M to binary values, the conditions in the multi-part cost function (5) can be rewritten and the optimisation problem

¹³In other words, we do not allow for splitting the execution between market and limit orders.

can be forked.

$$\text{cost}(M, \Delta) = \begin{cases} \text{imp.sh}(M, \Delta) & M = 1 \text{ and } \text{mk.imp}(V) \leq \sigma_{is} - BM_{is}, \\ & M = 0 \text{ and } \Delta \geq BM_{is} - \sigma_{is} \\ \beta \text{ imp.sh}(M, \Delta)^2 / \sigma_{is} & \text{elsewhere} \end{cases} \quad (12)$$

The selection decision between a market or a limit order placed at relative distance Δ^* becomes equivalent to choosing the minimum of the following three branches: $f(M = 1)$, $f(\Delta^* \geq BM_{is} - \sigma_{is} | M = 0)$ and $f(\Delta^* < BM_{is} - \sigma_{is} | M = 0)$.

I. When $M = 1$ (market order, negative price change relative to BM_{is}), one can discriminate between two cases – inside or outside the volatility bands – depending on the market impact size:

$$f(M = 1) = \begin{cases} BM_{is} + \text{mk.imp}(V) & \text{for } \text{mk.imp}(V) \leq \sigma_{is} - BM_{is} \\ \beta (BM_{is} + \text{mk.imp}(V))^2 / \sigma_{is} & \text{otherwise} \end{cases} \quad (13)$$

II. When $M = 0$ (limit order) and $\Delta \geq BM_{is} - \sigma_{is}$ (positive price change or negative price change smaller than the volatility threshold, i.e. inside the bands – identified in figure 4 with Δ^{II+} and Δ^{II-} , respectively). Let $A = \lambda_u \text{flow}(OBI) \text{size}(V)$. It follows that:

$$\Rightarrow f(\Delta | M = 0, \Delta \geq BM_{is} - \sigma_{is}) = BM_{is} - \Delta + A(\alpha_0 + \alpha_1 \text{dyn}(\Delta) + \alpha_2 \text{queue}(\Delta)) \quad (14)$$

III. When $M = 0$ (limit order) and $\Delta < BM_{is} - \sigma_{is}$ (negative price change larger than the volatility threshold, i.e. outside the bands – Δ^{III} in figure 4). As $\Delta > 0$, the precondition $BM_{is} - \sigma_{is} > 0$ must apply.

$$\Rightarrow f(\Delta | M = 0, \Delta < BM_{is} - \sigma_{is}) = \beta \frac{(BM_{is} - \Delta)^2}{\sigma_{is}} + A(\alpha_0 + \alpha_1 \text{dyn}(\Delta) + \alpha_2 \text{queue}(\Delta)) \quad (15)$$




Figure 4: Volatility bands for sell orders

In a simulation framework, the non-execution risk function can be based on the effective order queue component, which takes into account the actual state of the order book. This queue function increases in steps at random values because of the probable book gaps and stochastic depth sizes at various book levels. Thus, the queue function is not derivable and a numerical procedure has to be implemented. We propose an iterative procedure where a trajectory of potential solutions Δ_i starting at 0 is evaluated step by step with respect to the objective of minimizing $f(\Delta_i|M = 0)$ and the best candidate Δ^* is stored. Finally, the fitness of the best candidate for a limit order $f(\Delta^*|M = 0)$ can be compared to the fitness of a market order $f(M = 1)$ and the appropriate order type can be chosen.

If the queue effect is temporarily ignored, i.e. $\alpha_2 = 0$, the functional forms of the cost and risk functions can be exploited with the goal of identifying a *stopping point* for the numerical procedure, reached where the derivative of the fitness function $f(\Delta|M = 0, \alpha_2 = 0)$ is zero.¹⁴ Pre-identifying this critical point also allows for considering a sparser search space, by jumping from one book level to the next – since the only inflexion point is at the end of the search interval, all intermediary potential Δ_i situated within the order-book gaps can be ignored. Thus, the stopping point corresponding to branch (14) is given by:

$$\frac{\partial f(\Delta|M = 0, \Delta \geq BM_{is} - \sigma_{is}, \alpha_2 = 0)}{\partial \Delta} = -1 + 2\alpha_1 A \frac{\Delta - BM_{dyn}}{\sigma_{dyn}} = 0 \quad (16)$$

$$\Rightarrow \Delta_{1, \Delta \geq BM_{is} - \sigma_{is}}^S = BM_{dyn} + \frac{\sigma_{dyn}}{2\alpha_1 A} \quad (17)$$

If $BM_{is} - \sigma_{is} > 0$, the solution corresponding to the third branch (15) is:

$$\frac{\partial f(\Delta|M = 0, \Delta < BM_{is} - \sigma_{is}, \alpha_2 = 0)}{\partial \Delta} = -2\beta \frac{BM_{is} - \Delta}{\sigma_{is}} + 2\alpha_1 A \frac{\Delta - BM_{dyn}}{\sigma_{dyn}} = 0$$

¹⁴From a graphical perspective, the slope of the cost function decreases for $0 < \Delta < BM_{is} - \sigma_{is}$ and equals the constant -1 for $\Delta \geq BM_{is} - \sigma_{is}$. On the other side, the slope of the risk function increases for $\Delta > 0$. The inflexion point is situated where the slope of the adjusted risk equals the absolute value of the cost function slope.

(18)

$$\Rightarrow \Delta_{2, \Delta < BM_{is} - \sigma_{is}}^S = \frac{\beta BM_{is} \sigma_{dyn} + A \alpha_1 BM_{dyn} \sigma_{is}}{\beta \sigma_{dyn} + \alpha_1 A \sigma_{is}} \quad (19)$$

3 Market design

Since the main focus of this paper is to analyse the market dynamics generated by order placement decisions, we simplify the core investment decision processes and adopt the zero-intelligence paradigm with this respect. Moreover, as results are mainly benchmarked to Cui and Brabazon (2012) (henceforth referred to as “CB model”), we try to keep as much as possible common to their design. Therefore, the “population” structure is minimalistic as in Cui and Brabazon (2012), comprising of two agents – a buyer and a seller – and one market maker. Time is considered to be discrete with a millisecond granularity, and a single trading session of 8.5 hours corresponds to 30,600,000 milliseconds. At each millisecond, one of the buyer or seller agents is picked to trade with probability $1/2$. Each agent can choose between three possible actions: (i) do nothing with probability λ_o , (ii) submit a market or a limit order with probability $\lambda_m + \lambda_l$, or (iii) cancel the oldest outstanding limit order with probability $\lambda_c = 1 - \lambda_o - \lambda_m - \lambda_l$. Order sizes are random draws from a log-normal distribution, with the associated generating function: $\exp(\mu_{\text{size}} + \sigma_{\text{size}} r_{\text{norm}})$, where μ_{size} and σ_{size} are the location and scale parameters, and r_{norm} is a standard normal deviate. Whenever a side of the order book is empty, the market maker intervenes by filling it with three random *off-spread limit orders*, with the relative price drawn from a distribution with the following random number generating function: $xmin_{\text{offspr}} (1 - r_{\text{unif}})^{-\frac{1}{1 - \beta_{\text{offspr}}}}$, where $xmin_{\text{offspr}}$ and β_{offspr} are parameters, and r_{unif} is a uniform deviate between zero and one.¹⁵ In contrast to the stochastic cancellation process of agents’ limit order, the market maker’s limit orders are set to expire in five minutes. This allows them to have a longer life span than agent’s limit orders, leading to the formation of *order book gaps*, i.e blocks of adjacent price levels with missing quotes.

The main difference between the CB model and our model is that we replace the random order placement based on statistical distributions with a *microtrading strategy* as described in section 2 (therefore, our model will be

¹⁵This functional form of the random number generator actually produces values distributed to $1/x$ where $x \sim$ power-law, because of the minus sign in front of the fraction exponent. As a consequence, $xmin$ acts as an upper bound for the generated sample. The resulting distribution is discrete because of the tick size and its shape is plotted in figure 6. For comparison, the derivation of a power-law random number generator is provided in appendix A.

referred to as “Micro model”). The Micro model adds an extra layer dealing with order execution, which is separated and independent from the investment decision. The inputs of this layer are the size and direction of the order, as well as agent preferences regarding trading urgency, benchmarks and volatility bands. The current market state – defined by order-book liquidity, short-term price volatility – is also taken into account. The optimised micro-trading decision consists in generating the ultimate order submitted to the trading-venue, which can take the form of a market or limit order.

Heterogeneity over the agents’ sense of urgency λ_u is introduced by drawing, for each new order, random values from a mixture of two normal distributions. The two modes correspond to two types of agents, i.e. a patient type with λ_u closer to zero and an impatient type with λ_u around one. An absolute value operator is applied over the random draw to ensure $\lambda_u \geq 0$. Additional heterogeneity is provided by the log-normal random order sizes and the various conditions reflected by the order book.

4 Experimental setup and results

Both the CB and the Micro model are implemented in the same software framework in order to keep differences to the minimum – actually, the micro-trading agent is an extension of the original CB trader which overwrites only the placement decision; everything else, e.g. matching engine, agent pooling, order cancellation, market making is kept unchanged.

The parameters of the CB implementation have the same values as in Cui and Brabazon (2012), which are originally computed from a dataset for Barclays Capital from London Stock Exchange. Even if we have tried to reproduce the CB model based on its description in Cui and Brabazon (2012), the results still differ to some extent, e.g. market impact averages are lower in our implementation, so we cannot claim that we actually benchmark to the model implemented in Cui and Brabazon (2012), but to a similar model using our own coding.

The Micro model maintains the same parametrisation as in Cui and Brabazon (2012), where it applies: $\lambda_o = 0.9847$, $\lambda_m + \lambda_l = 0.008$, $\lambda_c = 0.0073$, $\mu_{\text{size}} = 8.2166$, $\sigma_{\text{size}} = 0.9545$, $xmin_{\text{offspr}} = 0.05$, $\beta_{\text{offspr}} = 1.7248$, initial mid-quote price 300.00, initial spread 0.50, tick size 0.01. The remaining parameters regarding the microtrading strategy are set in order to replicate the same order type frequencies as the ones generated by the CB model: 4% market orders, 10% in-spread and 86% off-spread limit orders. Also we have tried to roughly reproduce the stylized facts discussed in the rest of this section, but no intensive or automated calibration has been pursued. The chosen set of parameters is: default average daily volume $ADV = 77\text{m}$, book depth levels $N = 3$, OBI base $\mu = 2$, size penalty exponent $\eta = 0.8$,

$\alpha_0 = 0.1$, $\alpha_1 = 0.5$, $\alpha_2 = 0.25$, $\beta = 2.5$. The parameters of the distribution mixture associated with the sense of urgency λ_u are: $\mu_1 = 0.4$, $\sigma_1 = 0.2$, $\mu_2 = 1.1$, $\sigma_2 = 0.2$, and the probability of drawing from the first normal is 40%. The two benchmarks $BM_{is} = BM_{dyn}$ are chosen as the exponential moving average of trading price $\bar{p}_t = 0.95\bar{p}_{t-1} + 0.05p_t$, because this indicator is more stable than the mid-price, given the low market liquidity. The index t corresponds to the trade price time series and thus \bar{p} is updated every new trade. Similar, the two volatility bands $\sigma_{is} = 7\bar{\sigma}$ from (5) and $\sigma_{dyn} = 7.5\bar{\sigma}$ from (7) are multipliers of the estimated time-varying instant standard deviation computed as $\bar{\sigma}_t = \sqrt{0.95\bar{\sigma}_{t-1}^2 + 0.05r_t^2}$, where r is the percentage return.

Both models are run for 30 artificial days, each day with a different random seed. At the beginning of every day, the model is warmed up for 3,600,000 milliseconds (1 hour). All data, except for the warm-up period, is aggregated in one dataset and several statistics as well as two stylised facts are investigated. The “stylized facts”, i.e. empirical regularities exhibited by a wide range of financial time series, are commonly used to validate ABM designs and parametrizations. A wide range of stylized facts, both for high-frequency and aggregated data, are described in the literature, e.g. Cont (2001, 2011), Daniel (2006), Pacurar (2008), Chen, Chang and Du (2012). The class of intraday stylized facts can be associated to transaction data, order book shape and order flow. In this paper, we have chosen one stylised fact related to limit orders which states that the distribution of the relative limit prices decays asymptotically as a power-law, and one associated with market orders which were found to generate a non-linear concave price impact function of trade size. No stylised facts related to return or order flow are selected since both models assume purely random investment decisions and inter-events durations.

4.1 Emergent properties and determinants of order placement

In this subsection a list of summary statistics related to trading events and market dynamics are presented (see table 1). Moreover, we test how the results are linked to some explanatory measures, as proposed in the microstructure literature and previously detailed in section 2.

Overall, the Micro model generates a lower number of trades, less trading volume and a larger average spread than the CB model. These measures have not been taken into consideration during the modeling and parametrisation process, and carry no special meaning in discriminating between the two models. Regarding market maker’s activity, there are fewer interventions in the Micro model, since an agent is not allowed to send a market order that would consume the entire book liquidity. In other words, if the order is greater than the available book depth, the agent decides for a limit order,

	CB model	Micro model
Trades analysis		
Average daily volume	76m	62m
Av. daily no. of trades	32,319	25,090
Av. daily no. of cumulated trades ^a	17,519	12,088
Av. cumulated trade return ^b	-2.40e-06	-1.30e-05
Cumulated return variance	7.43e-07	2.15e-03
Quote analysis		
Average spread	0.11	0.77
Average percentage spread	0.04%	0.26%
Order types – effective frequencies		
Market maker orders	0.77%	0.09%
Agent market orders	3.75%	4.93%
Agent crossing limit orders	3.41%	0.003%
Agent in-spread limit orders	3.23%	10.23%
Agent spread limit orders	32.90%	8.17%
Agent off-spread limit orders	56.71%	76.67%

^aTrades initiated by the same order are merged as they belong to the same fill with multiple counter-party limit orders

^bThe resulting price of the cumulated trade is the weighted average price

Table 1: Summary statistics

no matter its aggressiveness. The market maker’s intervention is required only at the beginning in order to initially fill the order book – three limit orders on each of the two order book sides – or if one side of the book becomes empty due to order cancellation. On the other side, in the CB model, the market maker intervenes also during trade execution when there is not enough liquidity to fill an incoming market order. Eventually, the default market maker spread of 0.5 seems to emerge as a significant level for both models.

Figure 5 presents the spread binned scatter plot for the first run/day alongside the spread histogram with the mean value over all 30 runs.¹⁶ We can observe that for the CB model the most common spread value is equal to 0.10, i.e. the minimum possible given the tick size. This explains also the differences between the agents’ trading intentions and the ultimate limit order classification, e.g. when the spread is equal to the tick size, an inside-spread limit order becomes due to rounding either a crossing or a spread limit order. Furthermore, the relative price drawn from the power-law distribution can get rounded to zero because of the tick size, and thus some of the off-spread limit orders actually become spread limit orders.¹⁷

¹⁶Observations above the 99th percentile are considered outliers and are discarded in order for the graphs to zoom better on the data – however no observation is discarded for analysis purposes. Moreover, the different shapes of the histogram plots are not due to scaling.

¹⁷3.41% effective crossing limit orders vs. the model setting of 0.30% are recorded, 3.23% instead of 9.70% in the case of in-spread limit orders, 32.90% instead of 17.00% for spread limit orders, and 56.71% instead of 72.00% for off-spread limit orders.

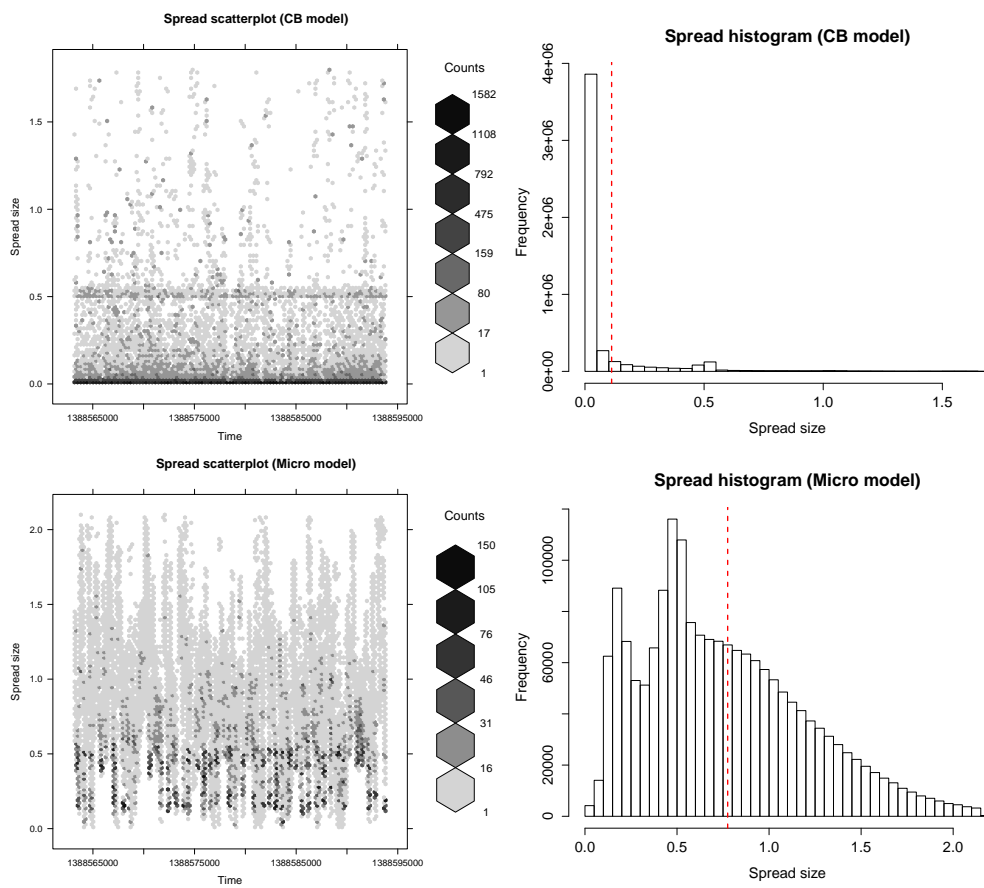


Figure 5: Bid-ask spread

In order to establish the relationship between the order type preference and various order and market factors, we will apply the logit regression in equation 20,¹⁸ where the dependent binary variable is the order type – market or limit – and the predictor variables are the indicator variable buy/sell, order size, market tightness, market impact, OBI and instant volatility.

$$\text{logit}(\text{Prob}(\text{isLimit} = 1)) = \beta_0 + \beta_1 \text{isBuy} + \beta_2 \text{size} + \beta_3 \text{spread} + \beta_4 \text{percMkImp} + \beta_5 \text{obi} + \beta_6 \text{vola} \quad (20)$$

The results are presented in table 2 and the coefficients can be interpreted as the linear impact on the log odds of choosing a limit over a market order (the reference group). The results are statistically significant due to the very large number of observations (7.34 million). Though, only some results are in accordance with the findings in the micro-structure literature – order aggressiveness is decreasing in the bid-ask spread, potential price impact of a market order and short-term volatility –, while others are contradictory – limit orders preference increases with order size, unfavourable overall supply-demand imbalance and sell orders. It is also to be noted that we have excluded important factors which we consider to be unobservable in real conditions, e.g. sense of urgency.

Coefficient	Estimate	Std. Error
(Intercept)	-0.016326*** ^a	0.004441
isBuy1	-0.117418***	0.003825
percSize	5.637183***	0.200678
percSpread	5.101229***	0.026782
percMkImp	0.130466***	0.012100
obi	0.377489***	0.003180
vola	45.851542***	0.103838

^a*** denotes statistical significance at < 0.001 percent level

Table 2: Logistic regression estimates

4.2 Relative price distribution of the off-spread limit orders

Zovko and Farmer (2002) define the relative limit price of a limit order δ as the difference between the limit price and the best quote on the same side of the market, i.e. the best bid for a buy order and the best ask for a

¹⁸Alternatively, we could differentiate between multiple levels of aggressiveness based also on the limit order price – market, cross, in-spread, spread, off-spread (can further be split in various intervals) – and use an ordinal logistic regression.

sell order. Furthermore, the standard procedure introduced in Zovko and Farmer (2002) takes into consideration only off-spread limit orders, i.e. only limit orders with positive relative prices ($\delta > 0$), while the rest – crossing, in-spread and spread – are discarded. The distribution of δ was found to decay asymptotically as a power-law, meaning that even if most of the limit orders are concentrated close to the best quotes, there are enough orders which are priced much less aggressively such that the distribution exhibits fat-tails. Different values for the characteristic exponent α associated with the power law probability distribution function $p(\delta) \sim \delta^{-\alpha}$ have been computed in the literature: Zovko and Farmer (2002) found $\alpha \sim 1.49$ for data from London Stock Exchange, while Bouchaud, Mézard and Potters (2002) and Potters and Bouchaud (2003) found $\alpha \sim 1.6$ for Euronext and NASDAQ.

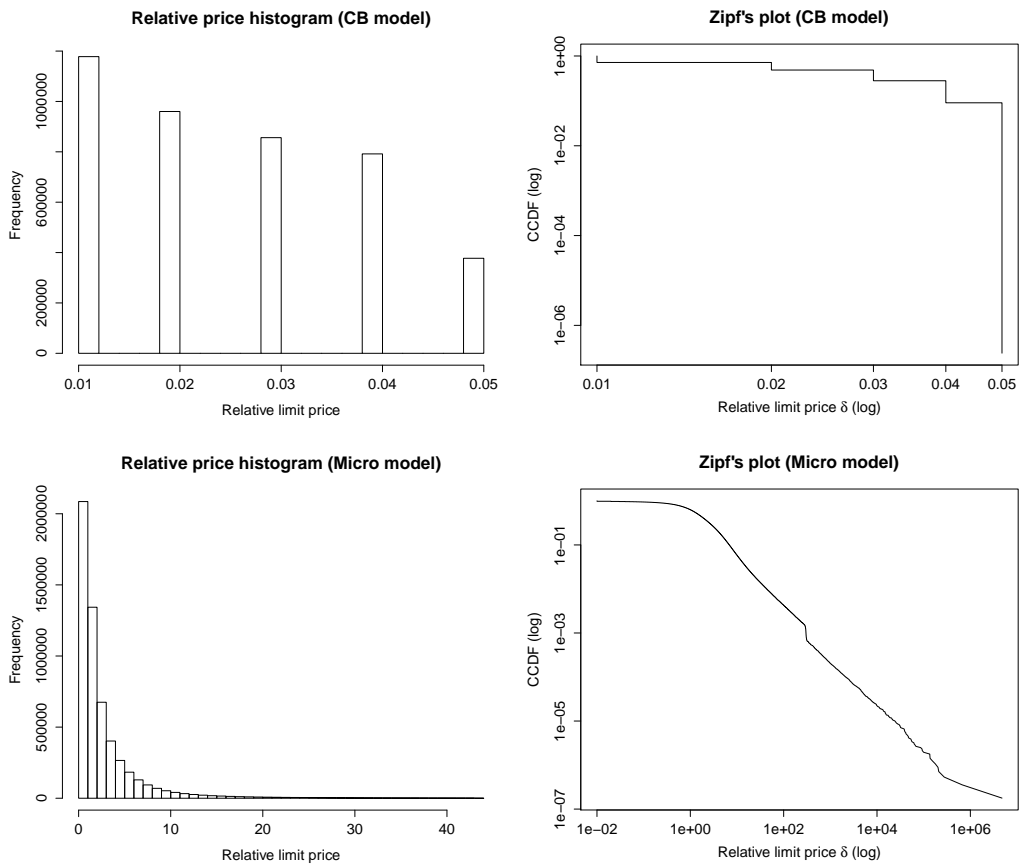


Figure 6: Relative limit price distribution of the off-spread limit orders

^a

^aThe bottom-left relative price histogram is zoomed by showing the x-axis up to the 99th percentile.

In the Micro model case, the power-law tail of the relative limit price distribution can graphically be assessed by the approximately linear shape of the Zipf plot on a log-log scale.¹⁹ We have also estimated the distribution expo-

¹⁹Plot the complementary cumulative distribution function of the ordered variable,

ment $\hat{\alpha} = 1.66$ for $xmin$ set to 0.5, using the R package ‘powerLaw’ which was developed based on the Santa Fe Institute recommendations.²⁰ A key role in replicating this stylised fact is played by the bimodal distribution of urgency coefficients λ_u , which allows for a number of passive limit orders with no chance of immediate execution to be sent to market by agents with low risk aversion. Otherwise, these passive orders could be the result of different agent’s evaluations or strategies, e.g. fundamentalists could place orders away from the current trading price, but no investment process is considered in the current model.

4.3 Market price impact

Market impact reflects the relationship between market order size and price impact, captured by the shift between the pre-trade and post-trade market equilibrium. Lillo, Farmer and Mantegna (2002, 2003) provide a method for computing the average market impact. Firstly, trades with the same time-stamp are aggregated and treated as a single transaction, as these are assumed to be part of a single market order which is matched against several outstanding limit orders. The market impact associated with each transaction is reflected by the difference in the logarithmic mid-quote price. Originally, the transaction size was measured in dollars, but we adopt the approach from Cui and Brabazon (2012) where the size of the market order is relative to the total daily trading volume. Finally, the data is divided into ten bins based on order size and the average price impact of each bin is computed. We also remove the upper outliers with respect to order size using a modified interquartile range rule,²¹ otherwise the number of observations in the higher bins would be very low if not zero.

A functional form of the market price impact is provided in the literature – Lillo et al. (2002, 2003) and Plerou, Gopikrishnan, Gabaix and Stanley (2002) consider the *empirical market impact function* for trade by trade data to be a power function of order size $\eta\nu^\gamma$, with exponent γ taking values between 0.2 and 0.6 for different stocks.²² This functional form is only an average property of the entire market. Since the order timing process is not observable and also not completely random – we can assume at least some intelligent trading taking into consideration available market liquidity – we

$\text{rev}(\text{cumsum}(1/N))$, where N is the total number of observation.

²⁰The estimation is highly sensitive to the value of the $xmin$ parameter. Thus, $x\hat{min}$ estimated using the package functions gives 314.38, for which the estimated $\hat{\alpha}$ is 2.0.

²¹The suspected outliers are all observations greater than the 99th percentile plus 1.5 times the interquartile range given by the difference between the third and the first quartile. The rule discards only 0.75% observations in the CB model and 0.65% in the Micro model.

²²Other studies have analysed the market impact on different time scales, by aggregating orders over time intervals of 5 or 15 minutes, as in Plerou et al. (2002) and Weber and Rosenow (2005), or by looking at the delayed price impact after 30 minutes as in Hopman (2007). The estimated exponent is slightly larger than in the case of tick by tick data and ranges from 0.33 to 0.75.

are confronted with an endogeneity issue which does not allow for the identification of the “true” relationship between order size and market impact only by analysing historical transaction data. Moreover, if the unconditional impact function would be concave, there would be no incentive to split a large order, as the total market impact of the smaller trades would be larger than the initial impact. Actually, Weber and Rosenow (2005) found that a *virtual price impact function* – computed by inverting the available order book depth as a function of return – is convex and is increasing faster than the concave average price impact function associated with effective market orders. Still, the overall average market impact represents a stylised fact which should emerge in a market with rational trading agents.

The average market impact computed on binned data – as in the standard procedure described previously – is depicted in the left plots of figure 7. In order to identify the relation with respect to the normalised order size, we estimate the coefficients of the power function $\beta_1 v^{\beta_2}$ using nonlinear least squares. The results presented in table 3 show a convex market impact function for the CB model, which is in line with the instantaneous price impact of Weber and Rosenow (2005) expected for randomly timed trading. On the other side, the Micro model function is concave with an exponent $\beta_2 = 0.67$, statistically significant different from zero and one.²³ The result is slightly larger than the empirical exponent levels found in the interval 0.2 - 0.6, and this can be explained by the lack of a thoroughly automated calibration, on one side, and by the major simplifications with respect to the investment process, on another side. Since the sample size for the binned data is very small, we repeat the estimations on all available data, but the results are not very different.²⁴ Both the individual as well as the average market impact are smaller in the Micro model, which can be explained by the selective trading strategy pursued by the agents aware of their market impact. This results is also in accordance with Weber and Rosenow (2005) who found that the virtual price impact is more than four times stronger than the actual one.

The right panels of figure 7 show a scatter plot (hexagonal binned) of the market impact for all transactions, as well as the fitted market impact function. In the CB model case, we observe that the data on the market impact axis is slightly bimodal, which raises the question of the validity of the mean estimate – even if the relative weight of the upper mode is very low. The mode around 0.001 is caused by the market maker’s intervention to prevent the order book from getting empty – even during executing of transactions – and resetting the bid-ask spread to its default value.²⁵

²³We have tested the restriction $\beta_2 = 1$ using the likelihood-ratio test.

²⁴We have also tried to discriminate between buy and sell initiated transactions, but found no difference.

²⁵The largest possible mid-quote difference after the intervention of a market maker is 0.27, determined by half of the spread expansion from a minimum of 0.01 to the market maker default 0.50 plus the range of the limit order power law distribution 0.05. When this shift is centred around the default price 300.00, the expected maximum market impact is

	CB model	Micro model
Market order percentage size		
min, max	6.98E-06, 1.14E-01	2.91E-05, 6.91E-02
50%Q, 75%Q	4.41E-03, 1.08E-02	4.51E-03, 9.39E-03
Bin size	0.0114	0.0069
Av. market impact range		
min, max	7.3E-06, 3.8E-04	7.3E-06, 5.5E-05
Power function estimates and standard errors for binned data		
β_1	6.54E-03 ^{***a} (1.72E-03)	3.43E-04 ^{***} (2.75E-05)
β_2	1.25E+00 ^{***} (1.06E-01)	6.69E-01 ^{***} (2.56E-02)
Market impact range		
min, max	0.00, 0.0103	0.00, 0.0014
Power function estimates and standard errors for all data		
β_1	5.92E-03 ^{***} (1.80E-04)	3.22E-04 ^{***} (7.21E-06)
β_2	1.21E+00 ^{***} (1.02E-02)	6.48E-01 ^{***} (5.30E-03)

^{***} and ^{**} denote statistical significance at < 0.001 and the 0.001 percent levels, respectively.

Table 3: Market impact measures

5 Conclusion and outlook

We have modeled the agent’s order placement decision as an optimisation problem which minimizes the risk adjusted execution cost, taking into consideration various micro-structure factors – such as order book liquidity, order flow proxied by the order book imbalance and transient volatility –, as well as intrinsic agents characteristics – such as the sense of urgency. We have derived an order submission strategy based on an iterative numerical procedure which allows for the efficient identification of the potential optimal limit price, taking into account the effective state of the order book. Next, we have integrated the order submission model into a zero-intelligence agent based model. Thus, we were able to assess the impact of replacing the original random order placement by the micro-trading strategy, with respect to two high-frequency stylised facts. Our model has successfully reproduced the power-law tail of the relative price distribution of off-spread limit orders, even if there is no explicit power-law component assumed and hard-coded into the agents’ design – thus it can be considered an emergent property. Regarding market orders, both the binned-average price impact as well as the individual price impact functions exhibit a concave shape – a trace of rational selective trading. On the opposite, in the absence of any intelligent trading decision, the expected market price impact shape is convex – confirmed by the results obtained with the alternative zero-intelligence agent based model.

Both the order submission model as well as the agent based model have some limitations which could be tackled in future implementations. The invest-

given by $\ln(300.14) - \ln(299.87) = 8.99\text{E-}4$.

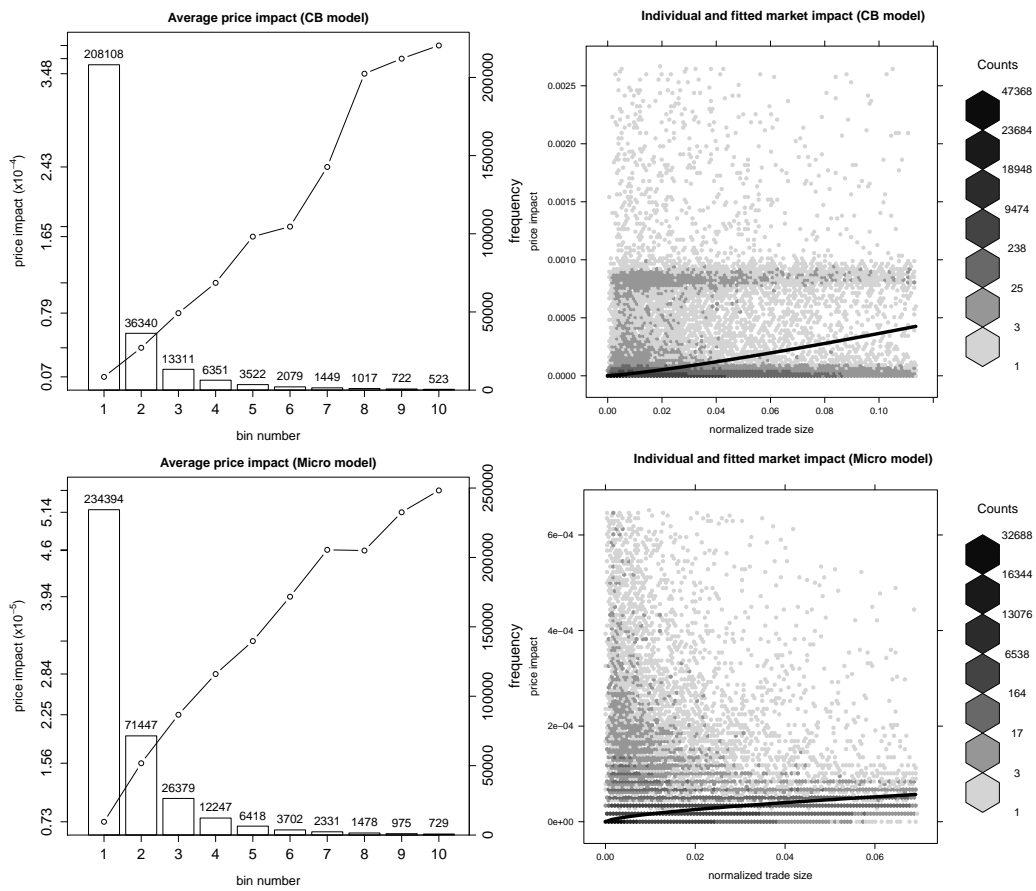


Figure 7: Market price impact

ment process is purely random and there is no relationship between agents' type and/or wealth with order sizes – usually volumes are correlated with strategy time-frame – or with risk aversion and trading urgency, respectively. The implementation of various missing model components could lead on the other side to relaxing some of the current model assumptions, e.g. the bimodal distribution of urgency coefficients could be replaced with a simpler distribution if heterogeneity is introduced into the model through modeling the investment decision. Moreover, no learning component is implemented. On one side, because of the short running time span of a single trading day, we can reasonably assume constant micro-trading strategies during the trading session. On the other side, agents do not adapt/react to the current market conditions in order to exploit them and, as consequence, market conditions do not change over time.

Several other stylised facts related to order book shape could be analysed, e.g. order book gaps. Also, even if the investment process is stochastic, some stylised facts related to price returns or order flow might emerge. More interesting would be the analysis of possible feedback loops related to price volatility and order book imbalance. Moreover, model calibration encompassing both the estimation of parameters as well as the choice of individual components, with the objective of better fitting various behavioural aspects, could be pursued.

The order placement model can be extended by taking into consideration other factors, e.g. intraday market trend, time of day, available trading time, prior order aggressiveness based on trading events clustering. Moreover, the market model can be enriched by adding other high-tech trading strategies such as algorithmic trading or high-frequency trading. We underline that, within the current model, the new computer-based strategies would interact with a market whose intraday dynamics is a little more similar to real markets than the dynamics generated by random placement decisions.

A Power-law random numbers generator

- Power-law pdf:

$$P(X = x) = f_X(x) = C x^{-\alpha}$$

- $\alpha = k + 1 > 1$ is the exponent of the power-law distribution, while k is known as the tail index (Pareto shape parameter)²⁶
- $C = (\alpha - 1) x_m^{\alpha-1}$ is a normalizing constant

$$\begin{aligned} F_X(x) &= \int_{x_m}^x f_x(x') dx' = \int_{x_m}^x C x'^{-\alpha} dx' \\ &= \frac{C}{-\alpha + 1} x'^{-\alpha+1} \Big|_{x_m}^x = \frac{(\alpha - 1) x_m^{\alpha-1}}{1 - \alpha} x'^{1-\alpha} \Big|_{x_m}^x = - \left(\frac{x_m}{x'} \right)^{\alpha-1} \Big|_{x_m}^x \\ &= - \left(\frac{x_m}{x} \right)^{\alpha-1} + \left(\frac{x_m}{x_m} \right)^{\alpha-1} = 1 - \left(\frac{x_m}{x} \right)^{\alpha-1} \end{aligned}$$

$$\int_{x_m}^{\infty} f_x(x') dx' = \int_{x_m}^{\infty} C x'^{-\alpha} dx' = - \left(\frac{x_m}{x'} \right)^{\alpha-1} \Big|_{x_m}^{\infty} = - \left(\frac{x_m}{\infty} \right)^{\alpha-1} + \left(\frac{x_m}{x_m} \right)^{\alpha-1} = 1$$

- Power-law random number generator

Let y a uniform distributed random variable in $[0, 1]$:

$$\begin{aligned} F_X(x) &= 1 - \frac{x_m^{\alpha-1}}{x^{1-\alpha}} = y \\ \Leftrightarrow x^{1-\alpha} &= \frac{1-y}{x_m^{\alpha-1}} \\ \Leftrightarrow x &= \left((1-y) x_m^{1-\alpha} \right)^{\frac{1}{1-\alpha}} = x_m (1-y)^{\frac{1}{1-\alpha}} \end{aligned}$$

$\Rightarrow F_X^{-1}(x) = x_m (1 - x)^{\frac{1}{1-\alpha}}$ is distributed as $f_X(x)$.

In the case of $x \in [x_m, x_M]$, x_m and x_M define the range of the distribution:

²⁶The correspondence is well explained in <http://www.hpl.hp.com/research/idl/papers/ranking/ranking.html>

$$\int_{x_m}^{x_M} C x'^{-\alpha} dx' = \frac{C}{-\alpha + 1} x'^{-\alpha+1} \Big|_{x_m}^{x_M} = 1$$

$$\Leftrightarrow C = \frac{1 - \alpha}{x_M^{1-\alpha} - x_m^{1-\alpha}}$$

$$\Rightarrow F_X(x) = \int_{x_m}^x C x'^{-\alpha} dx' = \frac{C}{-\alpha + 1} x'^{-\alpha+1} \Big|_{x_m}^x = \frac{x^{1-\alpha} - x_m^{1-\alpha}}{x_M^{1-\alpha} - x_m^{1-\alpha}} = y$$

$$\Leftrightarrow x = \left[(x_M^{1-\alpha} - x_m^{1-\alpha}) y + x_m^{1-\alpha} \right]^{\frac{1}{1-\alpha}}$$

$$\Rightarrow F_X^{-1}(x) = \left[(x_M^{1-\alpha} - x_m^{1-\alpha}) x + x_m^{1-\alpha} \right]^{\frac{1}{1-\alpha}}$$

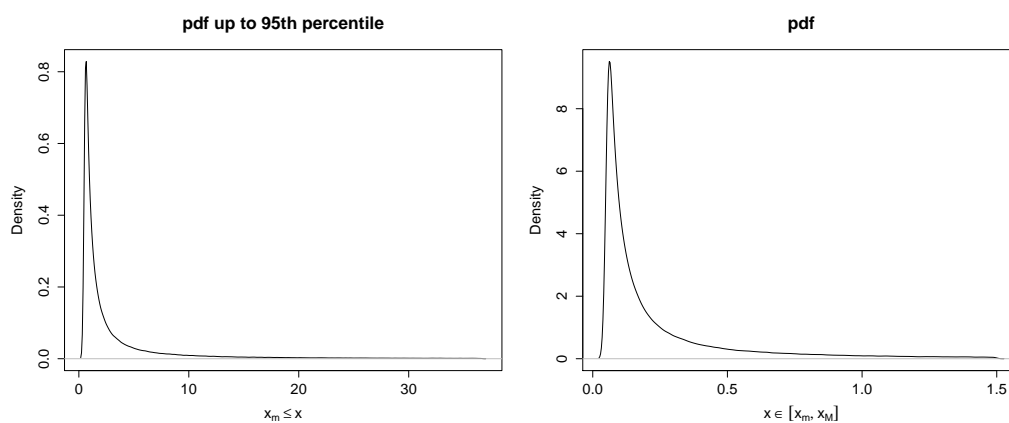


Figure 8: Power law distributions by simulation: lower-bounded $x_m = 0.5$, $\alpha = 1.7$ (left); lower- and upper-bounded $x_m = 0.05$, $x_M = 1.5$, $\alpha = 1.4$ (right)

References

- Arthur, W., Holland, J., LeBaron, B., Palmer, R. and Tayler, P. (1996). Asset pricing under endogenous expectations in an artificial stock market, *Available at SSRN 2252* .
- Beber, A. and Caglio, C. (2005). Order submission strategies and information: Empirical evidence from the nyse, *Technical Report 146*, International Center for Financial Asset Management and Engineering.
- Bollinger, J. (2001). *Bollinger on Bollinger bands*, McGraw Hill Professional.
- Bouchaud, J., Mézard, M. and Potters, M. (2002). Statistical properties of stock order books: empirical results and models, *Quantitative Finance* **2**(4): 251–256.
- Brock, W. and Hommes, C. (1997). A rational route to randomness, *Econometrica: Journal of the Econometric Society* pp. 1059–1095.
- Cao, C., Hansch, O. and Wang, X. (2009). The information content of an open limit-order book, *Journal of futures markets* **29**(1): 16–41.
- Chen, S.-H., Chang, C.-L. and Du, Y.-R. (2012). Agent-based economic models and econometrics, *Knowledge Engineering Review* **27**(2): 187–219.
- Chen, S.-H. and Yeh, C.-H. (2001). Evolving traders and the business school with genetic programming: A new architecture of the agent-based artificial stock market, *Journal of Economic Dynamics and Control* **25**(3): 363–393.
- Chen, S.-H. and Yeh, C.-H. (2002). On the emergent properties of artificial stock markets: the efficient market hypothesis and the rational expectations hypothesis, *Journal of Economic Behavior & Organization* **49**(2): 217–239.
- Chiarella, C., Iori, G. and Perelló, J. (2009). The impact of heterogeneous trading rules on the limit order book and order flows, *Journal of Economic Dynamics and Control* **33**(3): 525–537.
- Clauset, A., Shalizi, C. R. and Newman, M. E. (2009). Power-law distributions in empirical data, *SIAM review* **51**(4): 661–703.
- Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues, *Quantitative Finance* **1**: 223–236.
- Cont, R. (2011). Statistical modeling of high-frequency financial data, *Signal Processing Magazine, IEEE* **28**(5): 16–25.
- Cont, R. and Kukanov, A. (2012). Optimal order placement in limit order markets, *Available at SSRN 2155218* .

- Cui, W. and Brabazon, A. (2012). An agent-based modeling approach to study price impact, *Computational Intelligence for Financial Engineering & Economics (CIFEr), 2012 IEEE Conference on [proceedings]*, IEEE Press, pp. 1–8.
- Daniel, G. (2006). *Asynchronous simulations of a limit order book*, PhD thesis, University of Manchester.
- Farmer, J. D. and Joshi, S. (2002). The price dynamics of common trading strategies, *Journal of Economic Behavior & Organization* **49**(2): 149–171.
- Farmer, J. D., Patelli, P. and Zovko, I. I. (2005). The predictive power of zero intelligence in financial markets, *Proceedings of the National Academy of Sciences of the United States of America* **102**(6): 2254–2259.
- Fischer, T. and Riedler, J. (Forthcoming). Prices, debt and market structure in an agent-based model of the financial market, *Journal of Economic Dynamics and Control*.
- Gode, D. and Sunder, S. (1993). Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality, *Journal of political economy* pp. 119–137.
- Hopman, C. (2007). Do supply and demand drive stock prices?, *Quantitative Finance* **7**(1): 37–53.
- Johnson, B. (2010). *Algorithmic Trading & DMA: An introduction to direct access trading strategies*, 4Myeloma Press, London.
- Kovaleva, P. and Iori, G. (2012). Optimal trading strategies in a limit order market with imperfect liquidity, *Technical report*.
- Lillo, F. (2007). Limit order placement as an utility maximization problem and the origin of power law distribution of limit order prices, *The European Physical Journal B* **55**(4): 453–459.
- Lillo, F., Farmer, J. and Mantegna, R. (2002). Single curve collapse of the price impact function for the new york stock exchange, *arXiv preprint cond-mat/0207428*.
- Lillo, F., Farmer, J. and Mantegna, R. (2003). Master curve for price-impact function, *Nature* **421**(129): 176–190.
- Lux, T. (1995). Herd behaviour, bubbles and crashes, *The economic journal* pp. 881–896.
- Lux, T. (1998). The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fat tails of return distributions, *Journal of Economic Behavior & Organization* **33**(2): 143–165.

- Lux, T. and Marchesi, M. (1999). Scaling and criticality in a stochastic multi-agent model of a financial market, *Nature* **397**(6719): 498–500.
- Lux, T. and Marchesi, M. (2000). Volatility clustering in financial markets: A microsimulation of interacting agents, *International Journal of Theoretical and Applied Finance* **3**(04): 675–702.
- Pacurar, M. (2008). Autoregressive conditional duration models in finance: a survey of the theoretical and empirical literature, *Journal of Economic Surveys* **22**(4): 711–751.
- Pascual, R. and Veredas, D. (2009). What pieces of limit order book information matter in explaining order choice by patient and impatient traders?, *Quantitative Finance* **9**(5): 527–545.
- Plerou, V., Gopikrishnan, P., Gabaix, X. and Stanley, H. E. (2002). Quantifying stock-price response to demand fluctuations, *Physical Review E* **66**(2): 027104.
- Potters, M. and Bouchaud, J. (2003). More statistical properties of order books and price impact, *Physica A: Statistical Mechanics and its Applications* **324**(1): 133–140.
- Weber, P. and Rosenow, B. (2005). Order book approach to price impact, *Quantitative Finance* **5**(4): 357–364.
- Westerhoff, F. (2009). Exchange rate dynamics: A nonlinear survey.
- Westerhoff, F. (2010). A simple agent-based financial market model: direct interactions and comparisons of trading profits, *Nonlinear Dynamics in Economics, Finance and Social Sciences* pp. 313–332.
- Zovko, I. and Farmer, J. (2002). The power of patience: a behavioural regularity in limit-order placement, *Quantitative Finance* **2**(5): 387–392.