## Ca' Foscari University of Venice

# Working Paper

## Dan Ladley and Paolo Pellizzari

## The simplicity of optimal trading in order book markets

# The simplicity of optimal trading in order book markets

## Dan Ladley

*Department of Economics, University of Leicester, Leicester, LE17RH, UK*

## Paolo Pellizzari

*Department of Economics, Ca' Foscari Univeristy, 30121 Venice, Italy*

First Draft: February 2014

**Abstract**

A trader's execution strategy has a large effect on his profits. Identifying an optimal strategy, however, is often frustrated by the complexity of market microstructure's. We analyse an order book based continuous double auction market under two different models of trader's behaviour. In the first case actions only depend on a linear combination of the best bid and ask. In the second model traders adopt the Markov perfect equilibrium strategies of the trading game. Both models are analytically intractable and so optimal strategies are identified by the use of numerical techniques. Using the Markov model we show that, beyond the best quotes, additional information has little effect on either the behaviour of traders or the dynamics of the market. The remarkable similarity of the results obtained by the linear model indicates that the optimal strategy may be reasonably approximated by a linear function. We conclude that whilst the order book market and strategy space of traders are potentially very large and complex, optimal strategies may be relatively simple and based on a minimal information set.

*Address for correspondence*:
**Paolo Pellizzari**
Department of Economics
Ca' Foscari University of Venice
Cannaregio 873, Fondamenta S.Giobbe
30121 Venezia - Italy
Phone: (++39) 041 2346924
Fax: (++39) 041 2349176
e-mail: paolop@unive.it

# 1   Introduction

How should a trader optimally execute a trade? As academic understanding of financial markets and the effect of their structure has grown this question has become more nuanced and sophisticated. In early models markets were assumed to have a single price and react smoothly to changes in demand. In this context the question of optimal trading was often one of timing - when should a trader trade. As these models became more sophisticated and market makers started to play a role, issues such as order splitting and information hiding came to the fore. More recently with the inclusion of architectures such as order books the question has acquired new facets - not just when should a trader trade but also at what price and with what tool. Should a trader trade now with a market order? This guarantees trade at a specified, but potentially inferior, price. Or should they post a limit order in the belief that prices will improve and that greater returns will be made? The ability of traders to select the best order may potentially have a large effect on their profits. The size of this effect is increasing as algorithmic trading aimed at picking off inefficient submission becomes more common. The trader's choice will be contingent on their own information but importantly also the state of their environment - the order book. How this information should be used and just which pieces are important, however, is an open question.

In this paper we investigate the important and effect of information on trading strategies and market dynamics. We draw conclusion from two models. The first permits continuous prices, i.e. there is no minimum tick size, and trading strategies are conditioned on the prices of the best bid and ask quotes. In this model strategies are optimised through the use of Evolution Strategies, an optimisation technique based on evolution and adaptation of the most profitable strategies. In the second model traders submit orders on a discrete grid of ticks. Strategies are identified via the algorithm of [Goettler et al., 2005] ensuring that they are optimal for the specified game. We find that the amount of information traders use in their strategies has little effect either on the dynamics of the market or on the behaviour of the traders either under the optimal strategies or the linear approximations. We conclude that optimal trading strategies in a microstructure context

may be simpler than believed and importantly may be characterised by a linear combination of the information available at the best quotes. Further we conclude from this that models of financial markets do not need to concern themselves with interpreting the full information set available to traders strategies. Indeed, restricting consideration to the best quotes has little effect on results.

The dynamics of order book markets constitute complex situations through which traders interact. Traders and academics, when analysing or modelling these markets, are both faced with the task of combining large amounts of information to find an optimal strategy. One reason for this is the complexity of the environment - the amount of information available to traders in the book. Even under a Markov assumption - that the entire payoff relevant history may be captured by the current state of the book the information available is very substantial. Order books typically constitute price grids. At each discrete price there may be any quantity available to buy or sell (under the constraint that the highest buy price must be less than the lowest sell price). As a result the size of the information space is potentially infinite. Some of this information is undoubtedly more important than other pieces. Orders far from the best prices are unlikely to result in trades and therefore are potentially less important. Their presence in the book, however, would have an effect on extreme price movements and therefore may not be ignored. As such, different pieces of information will be more or less important and may have a smaller or larger effect on trading behaviour. Constructing the strategy - the optimal mapping between states and actions - in these markets is therefore a daunting task.

Since the relatively early stages of the academic microstructure literature models have been constructed in an attempt to do this. Frequently, however, this requires strong assumptions in order to maintain analytical tractability. For example [Parlour, 1998] considers a book of only 4 ticks in which two have infinite liquidity whilst [Rosu, 2009] assumes continuous prices and time, permitting instantaneous revision of quotes. There have also been attempts to model these markets and trading strategies numerically. The simplest case is the literature on zero intelligence models, e.g. [Ladley and Schenk-Hoppé, 2009], in which traders ignore information about the book and remove strategic considerations all

together. Whilst these models allow the full market architecture and realistic order submissions, they completely abstract from the central problem we are concerned with here. Other models such as [Chiarella et al., 2012] and [Chiarella et al., 2009] use exogenously specified rules for determining the choice between market and limit order submission and the appropriate price and quantity. These decisions are dependent on the traders demand and the best quotes in the market. They are restricted, however, by the pre-specified functional structure - there is no guarantee (or claim) that these strategies are optimal in this setting. A third avenue of research of particular interest is in the papers of [Goettler et al., 2005] and [Goettler et al., 2009], which use the numerical technique of [Pakes and McGuire, 2001] to solve an order book market game for a Markov perfect equilibrium in which the trading strategies are optimal. Whilst this may appear to solve the problem these techniques are still numerical demanding. The algorithm attempts to identify the optimal response in all relevant states in the state space. As the size of the order book (the number of orders present) increase, however, this state space grows exponentially. As a result this algorithm is only able to find optimal strategies under a constrained space - either information must be discarded or this algorithm is restricted to books with a relatively small number of ticks and with few orders present.

An important insight to this question is made by [Bouchaud et al., 2009]. In this paper the authors discuss how there may frequently be gaps in the order book - prices at which no orders are present. Even with these *static* gaps the book may be considered to be *dynamically* complete, i.e. orders will appear and accumulate as they are needed - they are issued on the fly to provide liquidity. As such, knowledge of many levels of the order book may not be fully revealing of the state of the world if there are traders present within the market that will provide liquidity when it is needed. Information beyond the best quotes may be unreliable. [Manahov et al., 2013] consider a related problem in which traders with different levels of cognitive abilities trade within financial markets. In this case cognitive ability is reflected by greater capacity for complex strategies and reasoning through larger genetic programs. They find that more intelligent traders enhance price discovery but damage price stability and liquidity. It is, however, important to emphasise

that this study is concerned with the cognitive ability of traders and not the information they have at hand or the size of the strategy space, as we focus on here.

As is the case in many other works, we assume traders are risk neutral profit-maximisers despite the fact that, as pointed out in [Parlour and Seppi, 2008], agents' decisions should in the end be coherent with their portfolio and consumption choices, which typically display risk-aversion. However, to keep the models numerically manageable, we use reduced-form trading preferences and assume that trading benefits, modelled through private values, proxy for the utility stemming from trading. See the first section of the extensive survey by Parlour and Seppi for a thorough analysis of this issue.

The paper is organised as follows. The next section gives details on the setup of the market, defines the strategies used by traders and the equilibrium concepts used in this paper. Section 3 describes the two models of optimal trading in a continuous double auction, based on the use of linear and Markov perfect equilibrium strategies. Simulation results are presented in Section 4, which reports aggregate data on the order book dynamic equilibria together with an illustration of the optimal strategy used by traders. Some discussion and conclusive remarks end the paper.

## 2  Setup

We model a standard order book based Continuous Double Auction (CDA) where at each time step a single trader enters the market. The trader is randomly allocated a type, buyer or seller with equal probability, and a positive reservation value $v \in V$ or positive cost $c \in C$ for a single unit of the traded asset. We assume that $V = C$ and $|V| = k$, i.e., that agents' values and costs belong to the same set of $k$ discrete positive values. Additionally, values and costs are uniformly drawn from $V$ and $C$ and are constant over time: $v_i \in V$ is the $i$-th buyer's private valuation of the asset and can be thought as the maximum price he will rationally pay for the asset. Symmetrically, $c_j \in C$ is the $j$-th seller's private cost for the asset and can be regarded as the minimum price at which he is rationally willing to sell the asset. We will assume, as done frequently in other works, that every agent buys or sells a *single* unit of the asset and, likewise, deal with cancellation

in a simplified and standard way: at the end of every time step each order stored in the book is cancelled[1] with (a small) exogenous probability $P_c > 0$ that is independent of time, state of the book and of the specific agent acting in that period.

At any time $t$ the book is a double sequence of outstanding unit orders

$$S_t = \{0 \le ... \le b_{3t} \le b_{2t} \le b_{1t} < a_{1t} \le a_{2t} \le a_{3t} \le ...\},$$

where $b_{1t}, b_{2t}, ...$ and $a_{1t}, a_{2t}, ...$ are the lists of buy and sell orders in the books. We often omit the time index for simplicity. The highest bid $b_1$ and lowest ask $a_1$ are referred as best bid and best ask, respectively. The distance $a_1 - b_1$ is referred to as the spread.

Traders submit a single order when they enter the market. The quantity is fixed at one unit but the trader must decide the price, computed using a function of the state of the book and their valuation: without loss of generality we describe the model for the $i$-th buyer (the situation for the sellers can be easily recovered, given the symmetry of the environment). The bidding function (or strategy)

$$B_{it} = f_i(a_{1t}, b_{1t}, I_{it})$$

provides the limit price $B_{it}$ (a *bid*, in this case), given the values of the best quotes $a_{1t}, b_{1t}$. The set $I_{it}$ contains all of the information available to the agent both public and private. This set may include the state of the book and their private valuation/cost. The submission of $B_{it}$ changes the book and results in an immediate trade, a *marketable order*, if the bid is greater than or equal to the best ask, i.e., $B_{it} \ge a_{1t}$. In this case, the two agents involved in the transaction get the associated profits: the buyer earns $v_i - a_{1t}$ and the $j$-th seller, who issued $a_{1t}$ previously, is paid $a_{1t} - c_j$ where $c_j$ is his cost. The book is then updated so that the best ask $a_{1,t+1}$ in the next tick will be given by $a_{2t}$. If instead $B_{it} < a_{1t}$, the new order is inserted[2] in the book, maintaining its ordering, to be possibly used in future trades. Any profit occurring after $t$ is accrued in the same way to the parties with no time-discount. In particular, if $b_{1t} < B_{1t} < a_{1t}$ the order is called *price*

---

[1] We never cancel the order in the time step in which it is submitted.

[2] We always use the standard price-time priority to break ties.

*improving* as it raises the current best bid. Bids for which $B_{1t} \leq b_{1t}$ are less aggressive as the relative limit price is queued after the best bid and, as a consequence, at least one trade is needed before execution is possible.

Notice that in this setup an immediate transaction may result from many different orders. Indeed, any bid for which $B_{1t} \geq a_{1t}$ produces a transaction and gives the very same profit, regardless of $B_{1t}$. In other words, there are non trivial subsets of bidding functions that are formally different and provide different limit prices but are *profit-equivalent*. This is especially true for strategies that often generate marketable orders and has implications for the interpretation of the numerical results of the following sections.

An equivalent description holds for the generic $j$-th seller whose limit ask is given by $C_{jt} = g_j(a_{1t}, b_{1t}, J_{jt})$, where $J_{jt}$ is the information set available (to the seller) at time $t$. We skip the details for brevity.

Agents are risk neutral and maximise the expected payoff (immediate or delayed), selecting a strategy to issue orders (bids or asks). Once the rules for the auction regarding cancellation and quantities, and the description of the agents are given, different models are obtained specifying the features of the strategies and the information that is processed. The $i$-th buyer will attempt to solve the problem

$$\max_{f_i \in \mathcal{F}} E[pay_{it} | O_{-i}, v_i], \tag{1}$$

where $pay_{it}$ is the random profit resulting from bidding what is prescribed by $f_i(a_{1t}, b_{1t}, I_{it})$ at time $t$, $\mathcal{F}$ is the set of admissible bidding function and $O_{-i}$ denotes the (fixed) strategies used by the other traders. To simplify notation, we omit $O_{-i}$ and $v_i$ whenever this is not harmful. More formally:

$$pay_{it} = \begin{cases} v_i - a_{1t} & \text{if the order is immediately executed: } B_{it} \geq a_{1t}; \\ v_i - B_{it} & \text{if the order is executed at some time } t' > t: B_{it} < a_{1t}; \\ 0 & \text{if the order is (randomly) cancelled before execution.} \end{cases}$$

The expectation in (1) is taken over all the states of the book that can be faced at

$t$ and over all the trajectories of states that can materialise for $t' > t$, starting from the initial condition $S_t$ at time $t$, under the use of strategies $O_{-i}$. Unless unrealistically strong assumptions are made, the previous optimisation problem is analytically intractable due to the path-dependency of the book and the intricacies of the auction mechanism. Finding a numerical solution of (1) is still a non trivial task. This, however, may be tackled in several ways, which will be detailed in what follows.

We will assume hereafter that all agents of the same type use the same strategy and are interested in the equilibrium situation in which no agent has the incentive to change strategy given what other agents do. In detail, we aim at approximately computing a set of bidding (asking) functions

$$O^* = \{f_1^*, ..., f_k^*, g_1^*, ..., g_k^*\}$$

such that for any buyer $i = 1, ..., k$, say, we have

$$E[pay_{it}|f_i^*, O_{-i}^*] \geq E[pay_{it}|f_i, O_{-i}^*], \forall f_i \in \mathcal{F}, f_i \neq f_i^*, \tag{2}$$

where $O_{-i}^*$ are the strategies optimally played by all the other agents/types. The intuition behind (2) is well known and requires an equilibrium to be a set of policies in which no agent has the incentive to deviate if the other traders stick to their optimal strategy.

## 3 The models

In this section, we describe two models of traders' behaviour in a CDA. Several features of the auction (most notably, due to the 'double' path dependency, uncertain execution and random cancellation) and the strategic interplay of different types make optimal decisions hard to select or even approximate.

The first model is arguably mimicking a minimal and memory less level of strategic reasoning. Traders submit their orders only based on the best quotes at the time of entering the market. Limit prices are simple weighted averages of $a_1$ and $b_1$ (plus a

constant). On the top of the best quotes, the information set available to any trader is the empty set. A similar model was used in [Pellizzari, 2011].

The second model, see [Pakes and McGuire, 2001, Goettler et al., 2005, Goettler et al., 2009], allows traders to make use of further information - the first $l$ quotes on either side of the book. The expected payoffs of all possible orders in each state of the book are explicitly computed by estimating the execution probability of each order submission (clearly, for marketable orders the execution probability is taken to be 1). As such the profit maximising order may be selected and, effectively, the price setting function may therefore be of arbitrary shape and complexity.

A more detailed description is given in the next subsections.

## 3.1 Linear strategies

We assume that the bid/ask to be submitted by traders at time $t$ is given by

$$B_{it} = f_i(a_{1t}, b_{1t}, I_i) = \min(\overline{B}, \alpha_i a_{1t} + \beta_i b_{1t} + \gamma_i), \tag{3}$$

for buyers and

$$A_{jt} = g_j(a_{1t}, b_{1t}, J_j) = \max(\overline{A}, \delta_i a_{1t} + \phi_i b_{1t} + \eta_i), \tag{4}$$

for sellers, where $\alpha_i, \beta_i, \gamma_i, \delta_j, \phi_j, \eta_j$ are real constant to be determined and $I_i = J_j = \emptyset, \forall i, j$. Essentially, all traders compute the limit price to submit by offsetting a linear combination of the best ask and the best bid. Slightly abusing terminology, we refer to these bidding functions as *linear strategies* in the following and notice that $f$ can be thought of as a function of the coefficients $\alpha, \beta, \gamma$ as well as a function of the best bid and ask. We enforce a minimal level of rationality and assume that no buyer bids more than some (large) constant price $\overline{B}$ and no seller's ask is satisfied with less than some (small) constant amount $\overline{A}$ but we do not otherwise constraint agents and they are free to pick any linear strategy even though, say, the resulting bid may exceed the private valuation of the asset and, hence, successful execution would cause a net loss. It is also clear from (3, 4) that bids and asks are continuous real values: this is to be contrasted with values

8

and costs that are discrete.

Using the previous linear formulation, we can describe the strategies of all traders as vectors in $\mathbf{R}^3$ so that the bidding function (3) for the $i$-th type is determined by $\mathbf{x}_i = (\alpha_i, \beta_i, \gamma_i)$. Analogously, the asking function for the $j$-th seller can be thought of as $\mathbf{y}_j = (\delta_i, \phi_i, \eta_i)$. Given a set of strategies for traders other than the $i$-th one:

$$O_{-i} = \{\mathbf{x}_1, ..., \mathbf{x}_{i-1}, \mathbf{x}_{i+1}, ..., \mathbf{x}_k, \mathbf{y}_1, ..., \mathbf{y}_k\},$$

he will attempt to maximise the profits solving the problem

$$\max_{\mathbf{x}_i \in \mathbf{R}^3} E[pay_i | \mathbf{x}_i, O_{-i}].$$

A trading equilibrium is a set of triplets (strategies)

$$O^* = \{\mathbf{x}_i^*, \mathbf{y}_j^*, i, j = 1, ..., k\}$$

such that

$$\mathbf{x}_i^* = \arg\max_{\mathbf{x} \in \mathbf{R}^3} E[pay_i | \mathbf{x}_i, O_{-i}^*],$$

for all buyers indexed by $i = 1, ..., k$ and with an analogous property holding for all sellers, $j = 1, ..., k$.

Numerically, the set of equilibrium strategies can be approximated by repeatedly solving the optimisation problem for each type, assuming all the other agents stick to their strategies, and running the algorithm over all types until "convergence is reached". The details of the method are outlined in [Pellizzari, 2011] and are based on Evolution Strategies. This optimisation technique, thoroughly surveyed in [Beyer and Schwefel, 2002], evolves the parameters of the population through a number of generations in which the tentative bidding functions are mutated, evaluated, deterministically ranked and discarded based on a fitness measure, before giving birth to the next generation. It is of particular interest here that a meta-parameter related to the strength of innovation is endogenously evolved together with the unknown parameters and can be used to gauge

whether convergence has been successfully reached.

## 3.2 Markov Perfect Equilibrium Strategies

The second model embodies a different approach in which beliefs of the probabilities of order execution are explicitly calculated. An equilibrium in this framework is a set of probabilities of execution for any limit order in any state of the book. Moreover, we require such an assignment $P$ of probabilities to be consistent, meaning that if agents trade based on the beliefs $P$, the realised probability of execution is indeed $P$, so that there is no discrepancy between beliefs and reality.

We assume that the bidding function $f_i$ takes values in $V$ and that the $l \geq 1$ best quotes are known on each side of the market[3]. We refer to $l$ as to the *information level* of the trader, with $l = 1$ being the situation in which no quotes other than the best bid and ask are known. More formally, the $i$-th buyer's bidding function is

$$f_i : V^{2l} \longrightarrow V, \quad (b_1, ..., b_l, a_1, ..., a_l) \mapsto B_{it},$$

where the bid $B_{it}$ maximises

$$P(B_{it}|S_t)pay_{it},$$

and $P(B_{it}|S_{it})$ is the (perceived) probability that the order will be executed in state $S_t$ either immediately or after some time. In equilibrium traders decide their bid based on the belief $P : V^{2l+1} \to [0,1]$ representing the probability that an order $B_{it} \in V$ issued in state $S_t \in V^{2l}$ at time $t$ will be executed (before exogenous cancellation).

The probabilities are iteratively found as outlined in [Pakes and McGuire, 2001], aiming at producing $P_n \to P$ for large $n$: for any bid $b \in V$ and a state $S$, at the start of the simulation we set $\forall b, S, \ P_0(b, S) = 1$ and $m_0^{b,S} = 1$.

It is important that the initial probability $P$ is optimistic to facilitate the exploration of the parameter space. The counter $m$ records the number of times a state has been

---

[3]We also consider a special case where $l = 0$. In this case prices are selected at random uniformly from the distribution $(0, v_i)$ for buyers and $(c_j, \overline{A})$ for sellers. This constitutes a Zero Intelligence (ZI) strategy as defined by [Gode and Sunder, 1993]

visited - here initialised to 1. The trader who arrives at the market in each period selects the optimal order based on the current estimates of probabilities. Each probability is updated each time step as follows. For a state in which an order executes:

$$P_{t+1}(b, S) = \frac{m_t^{b,S}}{m_t^{b,S}+1} P_t(b, S) + \frac{1}{m_t^{b,S}+1}, \quad m_{t+1}^{b,S} = m_t^{b,S} + 1.$$

For a state in which the order is cancelled: $P_{t+1}(b, S) = \frac{m_t^{b,S}}{m_t^{b,S}+1} P_t(b, S), \quad m_{t+1}^{b,S} = m_t^{b,S}.$

For states in which an order is neither cancelled nor executed: $P_{t+1}(b, S) = P_t(b, S)$, $m_{t+1}^{b,S} = m_t^{b,S}.$

A number of algorithmic devices are used to improve speed and avoid premature convergence.[4]

After running the model for $T$ time steps we test for convergence in probabilities. The model is run for a further $X$ time steps during which the updating procedure described above is not applied and probabilities are held constant. Through out this period the number of times orders are submitted in each state and the number of times those orders are executed are both recorded. At the end of the period for any state in which more than 100 orders are submitted the realised probability of order execution is compared with $P(b, S)$, namely the probability of execution estimated by the numerical algorithm. The average mean squared error over all such states is calculated. If this value is less than 0.001 the model is said to be converged, i.e., the equilibrium has been identified. If this is not the case the model is run for a further $T$ time steps with probability updating and the model retested. This is repeated until the model is converged. Once this is achieved statistics are collected from the model.

## 3.3   Further comments

The two models reviewed in the previous section have some similarities but are also different in important aspects. Agents in both frameworks share a common set of discrete values/costs and attempt to maximise the gain from trade in a risk-neutral fashion. In the

---

[4]Every 100,000 time steps we set $m_t^{b,S} = 1, \forall b, S$. Moreover, with probability $p_R$ rather than submitting the utility maximising order a trader instead submits a randomly chosen order in the current configuration. The effect of this is to help prevent local equilibrium. In particular due to poor early performance certain actions may no longer be chosen, however, as strategies are refined over time these orders may be once again acceptable. The random selection of these orders allows them to be reintroduced to the strategy.

Markov Perfect Equilibrium model, traders must pick a bid/ask among $k$ possible prices (ticks), explicitly computing the expected profit of each option available. The bidding function takes discrete values but is not restricted in any other way and, in particular, has the potential to reveal that optimal trading may be characterised by some form of non-linearity.

In contrast, agents using linear strategies can submit orders at any price and this model is not endowed, as was the case for the Markov Perfect Equilibrium market, with a natural tick-size. Hence, in the linear strategy equilibrium, the best quotes can be arbitrarily close at times and this can possibly increase the liquidity and efficiency of the trading process. The strategy of each type of buyer/seller is relatively simple and depends only on three coefficients, whereas a full set of probabilities must be known to take any trading decision in the other model. Importantly, the form of the bidding functions in the linear strategies market is rather restrictive and the possibility to devise or approximate any nonlinear trading scheme is ruled out. The following section presents the results of a set of numerical simulations and discusses the extent to which the differences between the two models have an effect on the book dynamics and traders' actions or profits.

In both models traders are risk neutral. If the traders were risk averse they would trade to minimise the risk of non-execution by placing fewer limit orders and more market orders. This may not necessarily result in a wider spread as, being risk averse, traders would place their orders less far back in the book. Hence, while the proportion of equilibrium orders may be different, the effects of information levels demonstrated in this paper are not likely to change.

# 4    Results

This section compares the book dynamics prevailing in equilibrium in the two strategic models. For comparison, we also provide results obtained in a market populated by non-strategic Zero Intelligence (ZI) traders.

Numerical results for the linear strategy model are based on 20 independent simulations and averages or other statistics are computed using an ensemble of 106,400 states

| Variable | Description | Value |
|---|---|---|
| $V$ | Buyer valuations | $\{0.05, 0.10, ..., 0.90, 0.95\}$ |
| $C$ | Seller Valuations | $\{0.05, 0.10, ..., 0.90, 0.95\}$ |
| $P_c$ | Probability of cancellation | 0.01 |
| $\overline{B}$ | Maximum Bid | 1.0 |
| $\overline{A}$ | Minimum Ask | 0.0 |
| $P_R$ | Probability to issue a random order | 0.01 |
| $X$ | Convergence assessment period | 1,000,000 |
| $T$ | Optimisation period | 1,000,000,000 |

Table 1: Values and description of the parameters used in the numerical simulations.

of the book.[5] For the Markov perfect equilibrium model results are calculated over 20 repetitions for each information level and are averaged over 1,000,000 states of the order book.

| Model | ZI | $l = 1$ | $l = 2$ | $l = 3$ | Linear |
|---|---|---|---|---|---|
| Best bid | 0.424 | 0.466 | 0.465 | 0.464 | 0.472 |
| Best ask | 0.576 | 0.534 | 0.534 | 0.536 | 0.524 |
| Spread | 0.152 | 0.068 | 0.069 | 0.071 | 0.051 |
| Quantity at best bid | 1.76 | 2.39 | 2.40 | 2.32 | - |
| Quantity at best ask | 1.76 | 2.42 | 2.44 | 2.40 | - |

Table 2: Summary statistics of the book under the four different information levels.

Table 2 shows the average state of the book under different models: together with ZI traders ($l = 0$), we have considered three different information levels $l$ and the use of linear strategies.

The market populated by ZI traders is substantially different from any strategic market, with much wider best quotes on average and an inflated spread. Clearly, the lack of strategic considerations in this case results in too many orders being randomly placed behind the best quotes and with a low probability of ending in a trade. Conversely, any market populated by strategic traders shows a much narrower spread, close to the gap between adjacent traders' values or costs. There is virtually no difference for different levels $l$ of information and little practical discrepancy between the set of the Markov equilibria and the linear strategies equilibrium. The average equilibrium spread using

---

[5]States are obtained from 20 independent simulations of 7 days of trading. We approximate a continuous flow of traders using a large population of 760 agents, 380 buyers and 380 sellers: hence, statistics are based on $106,400 = 20 \times 7 \times 760$ states.

linear strategies 0.051 compared to about 0.070 for the other models (regardless of $l$) but it must be noticed that in the latter cases the spread cannot be less than 0.05, as offers on opposite sides are discrete and cannot overlap.[6] As such the presence of a minimum price increment (tick) in the Markov model has only a small effect on the equilibrium market behaviour.[7]

To understand why the information level has little effect on behaviour it is beneficial to consider the problem faced by traders. In the model, in equilibrium, the traders' estimates of the probabilities of orders executing are always correct. For a given state $X$ in information level $l$ this probability is the average, weighted by frequency of occurrence, of all states that in information level $l + 1$ would map to state $X$. For instance, consider the state $X$ for $l = 1$ of $\{B_1 = 0.4, A_1 = 0.6\}$ (i.e., the best bid is 0.4 and the best ask 0.6). There are a large number of states in $l = 2$ which map to this, including $\{B_1 = 0.4, A_1 = 0.6, B_2 = 0.3, A_2 = 0.7\}$, $\{B_1 = 0.4, A_1 = 0.6, B_2 = 0.3, A_2 = 0.8\}$, $\{B_1 = 0.4, A_1 = 0.6, B_2 = 0.3, A_2 = 0.9\}$ etc. All of these states in $l = 2$ would be represented by $X$ in $l = 1$. The greater number of states allows traders to specify their strategy more finely but they do not measure the probability of execution over the set any more accurately. As such, there may be some states where traders are more aggressive at $l = 2$ than they would be in $X$ at $l = 1$ and, similarly, some where they are less aggressive. The chosen action at level $l = 1$ may therefore be viewed as the payoff maximising action averaged over all possible "equivalent" states at $l = 2$. This explains why the information level has little influence on the aggregate behaviour being actions averaged across all states.

A snapshot of the best quotes realised with linear strategies is depicted in Figure 1. The graph shows that there is considerable variability in the trading session as well as frequent periods in which the spread falls to minute levels (periods when the two lines

---

[6]The quantities at the best quotes for the linear model are not given as with continuous pricing there is never more than one order at this price.

[7]The effect of the width of the price grid - the number of ticks in the market - was also considered. Doubling the number of ticks in the price grid led to an increase in the spread of 50% whilst the quantities at the best quotes were found to be 50% greater under the smaller set of prices. Importantly, however, a larger price grid was found to have no effect on the behaviour of the model across information levels, i.e. for all information levels the spread and quantities available were the same.

Figure 1: Example of time evolution of best bid and ask in equilibrium using linear strategies. Best bid is given as a dashed line whilst the best ask is the solid line. Each time step corresponds to a single trader entering the market.

nearly intersect). This demonstrates why the average spread in the presence of linear strategies is smaller than in the Markov perfect equilibria.

Table 3 shows the distribution of spreads for all the markets. Again the statistics for the four markets with strategic traders are very similar. In all cases over half of the time the best bid and ask are within one tick of the equilibrium price. In 90% or more of the cases the spread is within two ticks and in nearly all cases the spread is within three ticks. In contrast, the ZI market shows much more variability in the spread. In only 14% of observations is the spread within one tick of the mid price indicating that the market is much more volatile and less efficient. This indicates that for markets populated by strategic traders the price is relatively stable and, importantly, there are only a small number of market situations which traders are faced with. As such the degree of strategic sophistication traders' require may be low.

Table 4 shows the relative shares of the type of orders submitted in different markets. Again, the ZI results differ markedly from the ones of the strategic models: marketable

|  |  | 0.45-0.55 | 0.40-0.60 | 0.35-0.65 |
|---|---|---|---|---|
|  | 0.45-0.55 | 0.14 | 0.28 | 0.39 |
| ZI | 0.40-0.60 | 0.28 | 0.49 | 0.65 |
|  | 0.35-0.65 | 0.39 | 0.65 | 0.85 |
|  | 0.45-0.55 | 0.61 | 0.79 | 0.80 |
| $l = 1$ | 0.40-0.60 | 0.78 | 0.97 | 0.98 |
|  | 0.35-0.65 | 0.80 | 0.98 | 1.00 |
|  | 0.45-0.55 | 0.62 | 0.79 | 0.81 |
| $l = 2$ | 0.40-0.60 | 0.78 | 0.95 | 0.97 |
|  | 0.35-0.65 | 0.81 | 0.97 | 0.99 |
|  | 0.45-0.55 | 0.62 | 0.78 | 0.81 |
| $l = 3$ | 0.40-0.60 | 0.77 | 0.94 | 0.96 |
|  | 0.35-0.65 | 0.80 | 0.96 | 0.99 |
|  | 0.45-0.55 | 0.50 | 0.69 | 0.74 |
| Linear | 0.40-0.60 | 0.68 | 0.89 | 0.95 |
|  | 0.35-0.65 | 0.72 | 0.94 | 0.99 |

Table 3: Distribution of ranges of bid and ask spreads for traders using ZI, Markov ($l = 1, 2, 3$) and linear strategic. Rows correspond to bid price and columns to ask prices.

| Model | ZI | $l = 1$ | $l = 2$ | $l = 3$ | Linear |
|---|---|---|---|---|---|
| Market Orders | 0.113 | 0.233 | 0.233 | 0.233 | 0.257 |
| Price Improving Limit Orders | 0.073 | 0.108 | 0.104 | 0.109 | 0.181 |
| Limit Orders at Best Quote | 0.045 | 0.162 | 0.167 | 0.161 | - |
| Limit Orders Behind Best Quote | 0.769 | 0.497 | 0.496 | 0.498 | 0.563 |

Table 4: Distribution of types of orders under the four different information levels along with number of cancellations and trades.

orders are halved with respect to the other markets, few orders are aggressively improving the extant quotes and, as a consequence, most of orders are placed behind the best quotes. These results broadly match those highlighted by [Ladley and Schenk-Hoppé, 2009] who found that the ZI model produced too few orders market orders and limit orders at the best quotes and too many behind the best quotes relative to empirical data. In reality, as well as in this model, strategic behaviour leads to fewer limit orders being wasted - being placed behind the best quotes with little chance of execution. Sophisticated traders choose not to submit these orders and submit price improving orders instead.

The market with linear strategies is slightly more efficient than the Markov markets, as seen in the fractions of market(able) orders, 25.7%, as compared to 23.3%. This implies that the traded volume is almost 5% bigger in the market with linear strategies than in the

Markov ones due to the smaller spread available in the first market. As before, orders at the best quote are meaningless in the linear model. We therefore, provide in the table only the share, 56.3%, of non-improving orders for the model with linear strategies. Despite some differences, all the strategic markets are rather similar as shown by a more accurate comparison, say, between the linear model and the one in which $l = 3$. The share 16.1% of "at the best quote" orders for the Markov model can be split in equal parts and tallied in the "improving" and "behind the quotes" orders, respectively, assuming that with equal probability an order at the best quote falls in either of the neighbouring category. In such a way the fifth column of Table 4 would have 18.9% of improving and 57.8% of "behind the best quote" orders, which should be compared to 18.1% and 56.3% of the sixth column, relative to the linear strategy equilibrium.

It is of interest to also look at the behaviour of the traders in equilibrium, particularly when they use linear strategies that are relatively simple. Recall that the models contemplate heterogeneous agents with different values and costs: while some may be strongly intra-marginal, feeling an intense pressure to finalise a trade to get profits, others - the extra-marginal ones - will basically have no chance to trade in equilibrium, being outstanding quotes at levels that do not make possible execution at a profit compared to reservation values. Moreover, as hinted in Section 2, even though different strategies are evolved in distinct simulations, they are however almost perfectly profit-equivalent.

A way to represent what agents do is to show what they bid/ask facing some frequently visited states of the book. We take the two symmetrical configurations in which the best quotes are 0.55, 0.50 and 0.50, 0.45, respectively. Figures 2 and 3 depict the median of the limit orders posted by intra-marginal sellers and buyers across all the simulations. When the best quotes are 0.50 and 0.45 (dashed in Figure 2), there is fierce competition among sellers who pushed the ask downwards to reach the equilibrium price. On the one hand, the strongest sellers, with costs equal to 0.05 or 0.10, issue marketable orders hitting the best bid and cashing 0.45 for one unit of the asset (see the black stars in the picture): they get less than the equilibrium price but trade is immediate and large profits are secured anyway. On the other hand, sellers whose cost exceeds 0.10 prefer to post limit orders

17

Figure 2: Trading behaviour of intra-marginal buyers and sellers facing best quotes 0.50 and 0.45. Black (red) stars denote market orders submitted by sellers (buyers) and black (red) solid lines show the median ask (bid) when limit orders are posted. The horizontal axis shows the costs for sellers and 1 less the values for buyers.

that are not immediately executed, see the black solid line: in particular, we observe that the median order is improving when $c = 0.15, ..., 0.35$ and behind the best quote when $c = 0.40, 0.45, 0.50$.

Buyers in Figure 2 find an attractive (low) ask and the ones whose value is larger or equal to 0.65 content themselves with a marketable order, see the red stars representing bids hitting the quote 0.50 and notice that the horizontal axis shows $1 - v$ for buyers. Agents with values $v = 0.60, 0.55, 0.50$ prefer to improve the outstanding best bid in order to gain priority, see the red solid line.

Figure 3 almost perfectly matches Figure 2, after swapping the roles of buyers and sellers. Even when the depicted behaviour is distinct, this results in minute differences in profits and even more so if one takes into account that the figures represent *median* behaviours. Take, for instance, the seller whose cost is 0.35 in Figure 3: he will decrease the ask to 0.502, virtually zeroing the spread and securing for himself an expected profit that is very similar to the one immediately cashed by the symmetric buyer whose value is 0.65 in Figure 2.

Figure 3: Trading behaviour of intra-marginal buyers and sellers facing best quotes 0.55 and 0.40. Black (red) stars denote market orders submitted by sellers (buyers) and black (red) solid lines show the median ask (bid) when limit orders are posted. The horizontal axis shows the costs for sellers and 1 less the values for buyers.

Overall, the pictures represents a rather sensible and, *ex post*, intuitive behaviour on the part of traders: strongly intra-marginal agents typically trade immediately using marketable orders, either because there is fierce competition on their side or because the quote on the opposite side is (already) captivating. The weakly intra-marginal traders improve the best quote to gain priority or patiently queue their orders in the hope that future, less unbalanced, states of the book will make their offers competitive.

# 5 Conclusion

In this paper we have used two models of order book markets to investigate the importance of information and strategic sophistication. The results provide insights into the effect and importance of information on optimal trading in order book based continuous double auctions. The statistical measures of market and trader behaviour differed little across levels of information. These statistics, however, were very different from those obtained

under the zero-intelligence model where lack of knowledge and the resulting random behaviour results in sub-optimal trading. We may therefore conclude that the crucial piece of information for traders in constructing their optimal strategy is knowledge of the best quotes. Further knowledge about the book conveys no value in this context: intuitively, this may be related to the dynamic nature of the book, where orders are likely to be added close to the best quotes as and if they are needed. As [Bouchaud et al., 2009] argue, the book may be dynamically complete even when quotes far from the best ones add little information or don't convey useful trading signals.

Key to the effect above is the finding that in equilibrium only a relatively small number of order book states occur as shown by the large percentage of observation in which the spread occupies a relatively narrow band around the equilibrium price. As such the possible situations that traders must develop optimal responses for are small in number. Traders strategies may therefore be relatively simple and easily learnt. Moreover, the similarity between the optimal Markov strategies and the linear approximation indicates that optimal trading may be approximated by a simple functional form further easing the cognitive burden placed on traders.

The work presented in this paper could be extended to consider more complex market settings. In this paper we have considered a relatively simple market - a fixed equilibrium price, unit quantities and exogenous cancellation. All three of these aspects could be made more sophisticated. A moving equilibrium price would exacerbate the risk for limit order submitters - increasing the chance of non-execution or picking off if the price moved away from or towards the order. Non-unit orders could increase the impact of a trader on the book as they would potentially be able to remove liquidity at multiple price ticks. Endogenous cancellation and resubmission of orders would allow traders to adapt their order placement to the changing state of the market. All three of these changes would possibly increase the value of information beyond the first tick. It was surprising in the current setting that only the first level of information was valuable. Identifying the requirements for this to be the case more generally, however, would be a potentially valuable advance.

# References

[Beyer and Schwefel, 2002] Beyer, H.-G. and Schwefel, H.-P. (2002). Evolution strategies: a comprehensive introduction. *Natural Computing*, 1:3–52.

[Bouchaud et al., 2009] Bouchaud, J.-P., Farmer, J. D., and Lillo, F. (2009). How markets slowly digest changes in supply and demand. In Hens, T. and Schenk-Hoppé, K. R., editors, *Handbook of Financial Markets: Dynamics and Evolution*, Handbooks in Finance, pages 57 – 160. North-Holland, San Diego.

[Chiarella et al., 2012] Chiarella, C., He, X.-Z., and Pellizzari, P. (2012). A dynamic analysis of the microstructure of the moving average rules in a double auction market. *Macroeconomic Dynamics*, 16:556–575.

[Chiarella et al., 2009] Chiarella, C., Iori, G., and Perello, 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.

[Gode and Sunder, 1993] Gode, D. K. and Sunder, S. (1993). Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–37.

[Goettler et al., 2005] Goettler, R. L., Parlour, C. A., and Rajan, U. (2005). Equilibrium in a dynamic limit order market. *Journal of Finance*, 60(5):2149–2192.

[Goettler et al., 2009] Goettler, R. L., Parlour, C. A., and Rajan, U. (2009). Informed traders and limit order markets. *Journal of Financial Economics*, 93(1):67–87.

[Ladley and Schenk-Hoppé, 2009] Ladley, D. and Schenk-Hoppé, K. R. (2009). Do stylised facts of order book markets need strategic behaviour? *Journal of Economic Dynamics and Control*, 33(4):817–831.

[Manahov et al., 2013] Manahov, V., Soufian, M., and Hudson, R. (2013). The implications of trader cognitive abilities on stock market properties. *Intelligent Systems in Accounting, Finance and Management*, page Forthcoming.

[Pakes and McGuire, 2001] Pakes, A. and McGuire, P. (2001). Stochastic algorithms, symmetric markov perfect equilibrium, and the 'curse' of dimensionality. *Econometrica*, 69(5):1261–81.

[Parlour, 1998] Parlour, C. A. (1998). Price dynamics in limit order markets. *Review of Financial Studies*, 11(4):789–816.

[Parlour and Seppi, 2008] Parlour, C. A. and Seppi, D. J. (2008). {CHAPTER} 3 - limit order markets: A survey. In Thakor, A. V. and Boot, A. W., editors, *Handbook of Financial Intermediation and Banking*, Handbooks in Finance, pages 63 – 96. Elsevier, San Diego.

[Pellizzari, 2011] Pellizzari, P. (2011). Optimal trading in a limit order book using linear strategies. Working Papers 16, Department of Economics, University of Venice "Ca' Foscari".

[Rosu, 2009] Rosu, I. (2009). A dynamic model of the limit order book. *Review of Financial Studies*, 22(11):4601–4641.