Time-dependent trading strategies in a continuous double auction

Shira Fano

*Department of Economics
Bocconi University, Milan*

Paolo Pellizzari

*Department of Economics
Ca' Foscari University of Venice*

First Draft: April 2011

Abstract

We model a continuous double auction with heterogeneous agents and compute approximate optimal trading strategies using evolution strategies. Agents privately know their values and costs and have a limited time to transact. We focus on equilibrium strategies that are developed taking into account the number of traders that submitted orders previously, as well as the number of who will submit subsequently. We find that it is optimal to place increasingly aggressive orders, according to a roughly linear schedule, and test the resulting equilibrium for robustness and accuracy.

Keywords

Continuous double auction, equilibrium trading strategies, evolution strategies.

JEL Codes

D53, D44, C61, C63

Address for correspondence:

Paolo Pellizzari

Department of Economics
Ca' Foscari University of Venice
Cannaregio 873, Fondamenta S.Giobbe
30121 Venezia - Italy
Phone: (+39) 041 234 ____
Fax: (+39) 041 2349176
e-mail: paolo@unive.it

This Working Paper is published under the auspices of the Department of Economics of the Ca' Foscari University of Venice. Opinions expressed herein are those of the authors and not those of the Department. The Working Paper series is designed to divulge preliminary or incomplete work, circulated to favour discussion and comments. Citation of this paper should consider its provisional character.

1 Introduction

The Continuous Double Auction (CDA) is one of the most commonly used trading platforms in financial markets. Trading optimally in such an environment is hard due to the sheer complexity of the decision options faced by traders, that choose quantity, type, limit price and timing of their orders based on past history and/or beliefs about others' behavior.

Given the analytical and computational problems involved in the analysis, models of CDAs usually make a number of simplifying assumptions, like unit-orders, no cancellation and little heterogeneity of traders. [Foucault, 1999], [Foucault et al., 2005] and [Rosu, 2009] solve analytically models of a CDA with a flow of patient/impatient traders, under several additional assumptions related to the exogenous arrival of market orders and the presence of a "trading crowd" that absorbs any quantity outside of a fixed spread. A robust feature of these models is that impatient traders always submit market orders whereas patient agents always prefer to issue limit orders.

We consider many types of buyers and sellers, with different evaluations and costs, respectively, that proxy varying degrees of impatience. We focus, in particular, on the time constraint faced by traders and assume that there is a finite amount of time to complete a transaction. Traders enter the market in random order and face the risk of missed execution that never occurs in the previously mentioned models. In equilibrium, they adapt their bids or asks to the number of traders that submitted orders previously, as well as the number of who will submit subsequently. This appears to be a rather realistic feature as presumably the trading behavior with hours to go is different from the one within few minutes from the closing bell. The seminal tournament among trading strategies in [Rust et al., 1994] suggested that "timing [...] is crucial for successful performance" and the two best programs, Kaplan and Ringuette, both took advantage of the current elapsed time. Interestingly, while simple timing information is useful, only a handful of programs tried more sophisticated approaches like predicting the responses of the opponents.

Our model has similarities with [Parlour, 1998] where a two-tick model of a CDA is presented. In Parlour's paper, heterogenous agents can submit a market order or queue at the fixed best bid B or best ask A . In the first case, they transact immediately at a less favorable price, while a limit order will be executed (at better terms) with some probability that endogenously depends on the accumulated orders at the quotes and on the behavior of the traders that are to appear before the end of the session.

We extend the previous model in that traders are allowed to develop a time-dependent submission strategy and can submit limit orders *at any price*. We do not explicitly model market orders but, instead, allow marketable

limit orders to be immediately executed. Hence, impatient agents seeking for immediacy can optimally post aggressive limit orders that will result in outcomes similar to the ones obtained by market orders.

To the best of our knowledge, such a model cannot be solved analytically and we resort to computational techniques to determine the optimal strategies for all traders (with different types). Formally speaking, we look for a Nash equilibrium of a game with stochastic payoffs, played by agents and by nature: the former pick a bid/ask when it is their turn to participate in the trading session; the latter samples the participants to the market and selects a queue that randomly affects the gains of the agents.

We compute the optimal actions using an evolutionary algorithm known as Evolution Strategies (ES) and find that, in equilibrium, agents submit more aggressive bids or asks as time elapses, in a roughly linear fashion.

The paper is organized as follows. The next section describes the CDA and the bidding functions used by traders, together with an illustration of ES. Our results are showed in Sect. 3 and the robustness of the findings is analyzed in Sect. 4, where we estimate the accuracy of the computed equilibrium strategies. We then conclude with some final remarks.

2 The model

Agents interact using a continuous double auction as a market platform. Traders on both sides of the market sequentially submit offers to buy and sell one item of a traded asset and can accept offers at any time during a trading session. Whenever a buyers' bid is greater or equal than a sellers' ask, the agents conclude a transaction; otherwise, their offers are stored in two separate limit order books and remain valid for future transactions.

In the *bid book*, offers are sorted from the highest to the lowest limit price. The order is reversed in the *ask book*, so that the most competitive offers, called *best bid* and *best ask*, are always the first in the respective side. Both books are reordered whenever a new offer arrives. When a transaction occurs, the offers involved in the exchange are removed. A standard rule of precedence decides the transaction price, namely the price is set by the quote that was submitted earlier.

The protocol can be implemented in different ways and real markets exhibit a large number of subtle variations. To simplify the model, as customarily done in the literature, we do not allow agents to cancel or resubmit orders during a single session and assume no informational asymmetries in that all traders access the same time-related information.

An equal number N buyers and N sellers are sampled from a set $\mathcal{P}$ of

traders. The i th buyer, $i = 1, \dots, N$ can redeem one unit of the exchanged asset for $v_i \in V$. Hence, his profit is $v_i - p$ if he trades at the price p . No profit is cashed in the case of no execution. Analogously, the j th seller has a cost $c_j \in C$. The profit is $p - c_j$ if a transaction occurs at price p and zero elsewhere. Buyers and sellers can't exchange roles and a buyer (seller) who acquires (transfers) a unit of the asset can't sell (buy) it later in the same session.

At the beginning of each day all agent are active, meaning they have the right to post offers. We allow traders to make only unit orders, so they have to choose prices, but not quantities. Agents have a unique chance to submit a limit order and, hence, they can transact at most once per day.

We assume that the order of arrival of traders is randomly chosen by nature, according to a uniform distribution over all possible queues. This reduces the complexity of the strategy space, eliminating issues related to the choice of the time to act and orders' cancellation.

The main source of uncertainty, in this framework, is the position $t \in \{1, \dots, 2N\}$ in the queue. Moreover, as the types of the agents (values or costs) are random, the profits of each agent depend, in general, on both his/her position in the queue *and* on the set of the other traders taking part in the auction.

Agents can be active for two reasons: either they are waiting for their turn to post an offer, or their offer is already on one of the books. The trading day ends when all $N + N$ agent have submitted their offer and the queue of traders waiting to place an order is exhausted. At the end of each day, all the remaining unfilled orders are cancelled and the books are erased. Differently from other models, traders have a limited time to transact, i.e., there are relatively few agents involved in each trading session. This can be interpreted as a thin market or, alternatively, as a device to model a short span of time in a longer day. In fact, realistically, at each point in time, only a handful of traders are competing on the books at viable prices. There is also a technical reason to restrict the number of traders: unless mechanisms to cancel orders are in place, offers will accumulate and the bid-ask spread will decrease to the point where there is a unique transaction price. In such a steady state, only intra-marginal agents trade, bidding and asking exactly the same equilibrium price. The resulting dynamics, like in [Dawid, 1999], is extremely flat when a large number of traders is considered.

A strategy (or bidding function) for a buyer with value v is the vector $w(v) = (b_1^v, b_2^v, \dots, b_{2N}^v) \in \mathbf{R}^{2N}$, whose t -th component is the bid to post when the agent has to act at position t in the queue. When it is not needed, we drop the value v (or c) and simply write w for a strategy. For instance, let the strategy used by a buyer be the vector $(0.4, 0.4, 0.4, 0.4, 0.4, 0.5, \dots, 0.5, 0.95, 0.95) \in$

$\mathbf{R}^{20}$. Then, the trader bids 0.4 if he acts as fifth or earlier in the queue; if not, he always submits 0.5, except if he is last or next to last in the queue, when he bids 0.95 (say, because the time is almost over.) We skip for brevity a detailed description of sellers' strategies, that clearly requires only trivial changes with respect to buyers.

Agents maximize the expected profit, that is dependent on their type (value or cost) and on their bidding function. Intuitively, a good strategy should be individually rational, that is, no bid should exceed the value of the trader. Moreover, it should strike the right balance in the so called immediacy-efficacy trade-off that is typical of a limit order market: the more aggressive the bid, the larger is the probability to trade but, conversely, the smaller is the gain.

Let $\pi(w_i, \mathbf{w}_{-i})$ be the expected profit of a buyer using the strategy w_i when all the other agents behave according to $\mathbf{w}_{-i}$. An equilibrium is a set $\mathcal{E} = \{w_i, i \in \mathcal{P}\}$ of strategies such that

$$\pi(u_i, \mathbf{w}_{-i}) \leq \pi(w_i, \mathbf{w}_{-i}),$$

for all i and all alternative strategies u_i . Loosely speaking, in an equilibrium, no agent has the incentive to change its strategy to increase his profits.

The determination of the equilibrium strategies is a hard computational task. Estimates of the expected profits can, to the best of our knowledge, be obtained resorting to simulation only. In other words, given a strategic profile $(w_i, \mathbf{w}_{-i})$, the quantity $\pi(w_i, \mathbf{w}_{-i})$ is an average over (many) sampled positions in the queue and over (many) sampled sets of trading partners.

2.1 Evolution strategies

Evolution strategies (ES) are population-based optimization algorithms. Superficially, they look similar to the well-known genetic algorithms (GA), sharing broad ideas related to evolutionary adaptation, selection, mutation and cross-over of the members of the population. A modern treatment of GA is in [Ashlock, 2006] and an insightful introduction to ES is provided in [Beyer and Schwefel, 2002]. ES are tailored to maximize real-valued functions of continuous variables, use deterministic selection in the creation of a new population and cleverly endogenize the mutation rate using specific evolvable meta-parameters in the chromosome. In contrast, the archetypical (and, perhaps, old-fashioned) GA uses binary encoding, stochastic selection and exogenously given mutation strength.

ES are conveniently described using the $(\mu/\rho \pm \lambda)$ -ES notation, to be read "mu slash rho plus or comma lambda". μ denotes the number of parents that

generate λ offsprings which are then evaluated and selected to form the next generation (again, with μ individuals). The value of λ usually exceeds μ to create selective pressure on the population. The recombination parameter $\rho \leq \mu$ refers to number of parents that are used to breed an offspring, so that, say, $\rho = 1$ means that descendants are clones of the (unique) parent and, at the other extreme, $\rho = \mu$ implies that an offspring is a function of all the μ parents. Finally, the “+” or “,” in the notation refer to the way deterministic selection based on fitness is performed: in comma-ES, the μ parents for the next generation are deterministically selected among the λ offsprings picking the ones with highest fitness. In other words, *all* the individuals of the previous generation are discarded. In plus-ES, instead, the μ parents of the next generation are deterministically selected in the wider set of $\mu + \lambda$ solutions formed by both (old) parents and (fresh) offsprings. Strong elitism¹ is clearly at work in a plus-ES, as the best solution is never discarded and, indeed, only improving offsprings are included in the next generation.

In more detail, we describe how our $(\mu/\rho, \lambda)$ -ES algorithm works for buyers having a specific value, assuming the strategies played by other agents are fixed in a given batch of s simulated sessions. Obvious modifications apply to sellers. In the following, z denotes a real standard normal random draw variable and $\mathbf{z}$ is a λ -dimensional standard multivariate normal vector. Each time z and $\mathbf{z}$ are used, they are resampled independently from the previous draws.

1. Set $g = 0$.

Initialize the population $\mathcal{A}^{(g)} = \{(w_i^{(g)}, \sigma_i^{(g)}), i = 1, \dots, \lambda\}$ of λ individuals ($w_i \in \mathbf{R}^{2N}, \sigma_i \in \mathbf{R}$). The vectors $w_i^{(g)}$ dictate the bid to issue for each of the $2N$ possible positions in the queue, while $\sigma_i^{(g)}$ is a meta-parameter related to mutation.

2. Repeat (“years” of s sessions, where $N + N$ agents trade in a CDA).

- (a) Compute average profit $\pi_i^{(g)}, i = 1, \dots, \lambda$, and let $\mathcal{F}$ be the set of indexes of the μ agents (out of λ) with the largest average profit in year g (comma-ES).

- (b) For $l = 1, \dots, \lambda$ do:

- i. Sample (without replacement) a random subset $\mathcal{F}' \subseteq \mathcal{F}$ with ρ elements.

¹*Elitism* refers to algorithms where the elite, namely some set of best solutions, is guaranteed to survive to the next generation.

Set

$$w = \frac{1}{\rho} \sum_{k \in \mathcal{F}'} w_k^{(g)}$$

$$\sigma = \frac{1}{\rho} \sum_{k \in \mathcal{F}'} \sigma_k^{(g)}$$

- ii. Set $\sigma_l^{(g+1)} = \sigma \exp(\tau z)$.
- iii. Set $w_l^{(g+1)} = w + \sigma_l^{(g+1)} \mathbf{z}$.
- (c) Set $\mathcal{A}^{(g+1)} = \{(w_l^{(g+1)}, \sigma_l^{(g+1)}), l = 1, \dots, \lambda\}$ (this is the next generation)
- (d) Set $g = g + 1$

3. Until termination ($g \leq 250$).

Step 1 initialize a random population of λ individuals at generation $g = 0$. Observe that each member of a population is a couple of one strategy vector $w_i^{(g)} \in \mathbf{R}^{2N}$, with tentative bids to be issued in position $1 \leq t \leq 2N$ in the queue, and one real-valued meta-parameter $\sigma_i^{(g)}$, that controls the size of the mutations of the components of $w_i^{(g)}$. Successful evolution will, at the same time, produce individuals with effective w_i (in terms of average profit) and low σ_i (to reduce the chance of catastrophic mutations that preclude survival in future generations).

Step 2 is the main loop in the ES algorithm, repeatedly evaluates the average profit in s sessions and select the best μ individuals to form the parent population $\mathcal{F}$ in 2a. Then, λ offsprings are created (2b) by randomly picking ρ parents, taking averages of both the sampled strategies and meta-parameters (i), mutating first σ (ii) and then w (iii). In 2c, the just generated recombined and mutated individuals form the next generation $\mathcal{A}$, and the index g is incremented (2d).

Step 3 check whether some (more or less arbitrary) terminal condition is satisfied. In this paper, we stopped the ES after 250 years of s days, i.e., trading sessions.

The vigilant reader may notice that our description slightly differs from the one given in Figure 1 of [Beyer and Schwefel, 2002], line 3, in that we start with a population of size λ to generate μ parents only *after* the average profits in a CDA are simulated. The two formulations are perfectly equivalent but allow us to keep constant the number of agents involved in the market. For the same reason, comma-ES were preferred to plus-ES, where, at times, $\mu + \lambda$ buyers with the same values would have to be taken into account. It is

Table 1: Parameters of the model. The upper part shows the choices relative to the agents, the lower part the ones related to ES.

Parameter	Value	Description
N	10	Number of buyers and sellers in a session
V	$\{0.05, 0.1, \dots, 0.9, 0.95\}$	Set of values for buyers
C	$\{0.05, 0.1, \dots, 0.9, 0.95\}$	Set of costs for sellers
$\mathcal{A}$		Entire population (20 agents for each $v \in V$ and $c \in C$)
s	200	Number of trading session in one generation
λ	20	Number of offsprings
μ	10	Number of parents
ρ	1,5,10	Recombination parameter
τ	$1/\sqrt{4N}$	Meta-mutation rate in ES

also noteworthy that all the (important) minutiae related to the trading in a CDA are under the hood in 2a, where we succinctly say that “average profits are computed”. In particular, profits in 2a are computed in trading sessions with buyers and sellers of *any* type, whereas evolution occurs always within individuals with same values or costs: s sessions are used to compute the profits of all groups; based on these fitness measures, parallel ES algorithms evolve the strategies of all the agents of each type; such strategies are used in the next s sessions and the whole process repeats.

If desired, different strategies adopted by agents with the same values or costs can be interpreted as competing trading rules used by the same physical agent in several sessions. The agents then learn the best rules by evolving the most effective behavior within the set of their own trading strategies.

3 Computational results

In this Section, we discuss our results for a representative set of parameters’ values. Table 1 summarizes our choices:

Each sessions involves $N = 10$ buyers and sellers. Hence, 20 agents are randomly selected and act in a session according to a random queue. As discussed before, a relatively low N must be used to avoid trivial dynamics and we used the intermediate size of the market described in [Fano et al., 2011].

A strategy w (bidding function) is then encoded in a vector with 20

components representing the bid/ask to post in each of the possible positions $t = 1, \dots, 20$ in the queue. The realized payoff is a random variable whose mean is $\pi(w_i, \mathbf{w}_{-i})$, estimated taking the average profit over $s = 200$ sessions.

The values and costs of traders are sampled in the discrete and equal sets V and C . A standard approach for the computation of equilibria would require to solve a global optimization problem with 760 real variables: in fact, each strategy requires 20 entries and there are $19 \times 2 = 38$ different types.

As far as the ES parameters are concerned, we pick $\lambda = 20$ offsprings and $\mu = 10$ parents. We tried different ρ s, to investigate the robustness of the search results with respect to different recombination mechanisms.

As outlined in Sect. 2, we decompose the task to find an equilibrium in 38 subproblems, iteratively maximizing the profits of each type in isolation over s sessions, keeping constant the actions of the other types. In such a way, we solve the 38 disjoint problems

$$\max_{w_i} \pi(w_i, \mathbf{w}_{-i}),$$

relative to agents having the same $v \in V$, where $\mathbf{w}_{-i}$ denotes the strategies played by all remaining agents.

Figure 1: Left panel: bunches of optimal bids for buyers with values 0.95, 0.75 (top) and equilibrium asks for sellers with costs 0.05, 0.25 (bottom). The results are relative to $\rho = 5$. Right panel: probabilities of immediate (solid), delayed (dashed) and missed execution (dash-dot) as function of the values of buyers. The results are relative to $\rho = 10$.

Figure 1, on the left, depicts a bunch of 20 equilibrium strategies (averaged over 50 runs of the ES) for some intramarginal buyers and sellers.

In equilibrium, buyers increase their bid as time elapses and sellers, symmetrically, lower their ask. At the beginning of the session, when only the first traders in the queue have posted bids or asks, the orders tend to cluster around 0.5. This is done as plenty of time is available and delayed execution is likely, due to high number of agents that are still to come. Close to the end of the session, however, orders are much more aggressive, trading off gains for a higher execution probability. We numerically computed the equilibrium strategies of all agents in 50 independent runs, for each $\rho = 1, 5, 10$, totalling 150 experiments. In all cases, it is optimal to place increasingly aggressive orders, according to a roughly linear schedule and, hence, we have reasons to think that this is a robust feature of the evolved equilibrium strategies.

The right part of Figure 1 graphically illustrates some interesting features of the equilibrium outcomes. We plot the probability of immediate execution (solid), of delayed execution through a limit order (dashed) and of missed execution (dash-dot) as functions of the value of the trader (a symmetric plot for sellers is skipped.) The solid line shows that the most impatient traders, with high values, issue market(able) orders that results in immediate execution more than 60% of times. The remaining orders are filled with some delay even in the case of deeply intramarginal buyers. This is a novel and, we believe, realistic aspect of the model that is not present in some of other models of CDAs. Most of deeply intramarginal buyers, say when $v \geq 0.8$, transact almost certainly in equilibrium.

Moderately intramarginal agents trade less frequently with market orders, very often experience delayed executions (with a peak when $v \approx 0.6$) and there is an increased probability of no trade. Marginal traders issue market orders very rarely but still have a significant probability to have their orders filled at a later time. As expected, very few transactions involve deeply extramarginal traders that have a probability of missed execution that approaches 100%.

Overall, this picture is consistent with a world where impatient traders often (but not always) ask for immediate liquidity that is often provided using limit orders by patient agents. Buyers and sellers of intermediate strength roughly obtain the same proportion of immediate and delayed executions.

4 Quality of the equilibrium and robustness test

We investigate in this Section the accuracy of the claim that the output previously described is an equilibrium set of strategies, focusing on the lack

of any incentive to profitably deviate.

To evaluate the quality of the equilibrium found with ES, we use the notion of ε -equilibrium. A set $\mathcal{E}$ of strategies is an ε -equilibrium if, for all players i , for every strategy w_i and for every deviation u_i we have

$$\pi_i(w_i, \mathbf{w}_{-i}) \geq \pi_i(u_i, \mathbf{w}_{-i}) - \varepsilon,$$

where $\varepsilon \geq 0$ and there is no $\varepsilon' > \varepsilon$ such that the previous equation holds. We can interpret ε as measure of the distance of $\mathcal{E}$ from a (true) Nash equilibrium, meaning that at least one agent can obtain an amount ε of additional profits by abandoning his equilibrium strategy w_i to switch to u_i . That is to say that an ε -equilibrium restricts the profit attainable using a strategic deviation to be smaller than ε and coincides with a Nash equilibrium only if $\varepsilon = 0$.

Given $\mathcal{E}$ we can, in principle, compute the profits' differentials $\pi_i(w_i, \mathbf{w}_{-i}) - \pi_i(u_i, \mathbf{w}_{-i})$ for all possible deviations, keeping constant the strategies of all the other players $\mathbf{w}_{-i}$.

However, checking for all realizable deviations of vectors in $\mathbf{R}^{20}$ is unquestionably too demanding and an exhaustive comparison would just be practically impossible. Therefore, to solve this problem, we somewhat arbitrarily decided to reduce the complexity of the task by considering a subsample of deviations for some specific agents. We are aware that, by definition, we will underestimate the real ε but, at the same time, we provide useful simulation-based measures of the accuracy reached by ES in approximating the equilibria.

Without loss of generality, we will focus in what follows only on intramarginal buyers, sampling agents with probability proportional to their value. Deep intramarginal traders, in fact, obtain higher gains from trade and, on the other hand, can experience greater losses if they are not using an optimal strategy. Thus, we expect them to potentially benefit from strategic deviations.

Regarding which deviations to analyze, we consider three alternative ways to sample candidates, leading to three different families of ε . First, as in [Fano et al., 2011], we examine simple constant strategies where the bid does not vary with the position in the queue.²

Second, we create a deviation by adding a random quantity to each of the components of an existing strategy $w \in \mathbf{R}^{20}$. We, therefore, consider the vector w' whose j th component is $w'_{(j)} = w_{(j)} + \tilde{u}_{(j)}$, where $\tilde{u}_{(j)}$ takes the values -0.05, 0, 0.05 with equal probability.

²In detail, we require the deviations to be individually rational and sample the constant from a uniform distribution in $[0.45, v]$, where v is the value of the agent that deviates. Bids smaller than 0.45 would remain unmatched with high probability and, hence, would surely produce tiny, if any, additional profit.

Third, taking into account that we have 50 runs for each parameters' setup, we randomly replace an equilibrium individual strategy in *one* run with a deviation from *another* run. The rationale for this procedure is that bids that were optimal in one environment may also be able to provide good performance in another instance of the problem.

Table 2: Estimates of ε and ε^w , obtained using different methods to sample deviations, for different values of ρ . The p -value of the test that ε or ε^w are equal to zero are also shown, together with the value v of the buyer that deviates from the equilibrium. All entries, as explained in the text, are based on 5000 simulated deviations.

		ε	p -value	Value v	ε^w	p -value	Value v
$\rho = 1$	1	0.0057	0.416	0.85	0.133	10^{-14}	0.95
	2	0.0048	0.083	0.90	0.066	0.0004	0.95
	3	0.0021	0.062	0.95	0.110	10^{-5}	0.85
$\rho = 5$	1	0.0034	0.091	0.95	0.036	0.0006	0.80
	2	0.0019	0.027	0.95	0.022	0.0023	0.65
	3	0.0018	0.051	0.90	0.037	0.0008	0.80
$\rho = 10$	1	0.0025	0.340	0.85	0.032	10^{-5}	0.60
	2	0.0029	0.092	0.95	0.016	0.0004	0.90
	3	0.0010	0.014	0.85	0.073	10^{-5}	0.75

For each run, we have computed by simulation the profits arising from 100 independent deviations for each of the three methods described above. Table 2 shows the results, as a function of the value of the recombination parameter $\rho = 1, 5, 10$. In this and in the following table, every row begins with 1, 2 or 3, depending on the way deviations are generated. As an example, the row labelled “2” for $\rho = 5$ demonstrates that an additional mean profit of 0.0019 can be gained by the strongest buyer whose type is 0.95. With a p -value of 0.027, this amount can be considered different from zero at the 5% but not at the 1% level of statistical significance.

A look at the left part of the table, reveals that the size of ε , regardless of the type of deviations, is rather small and of the order of 10^{-3} . This means that, on average, minute profits can be obtained diverging from the equilibrium. Statistically, the amounts are often significant at the 10% level, but never at the 1% significance level and, in general, the values are larger for deviations of the first type and smaller when generated using the third method. Moreover, it is apparent from the sizes of the various ε s that the equilibria computed using $\rho = 1$ are less accurate than the one obtained with $\rho = 5$ or $\rho = 10$. Indeed, deviations from the equilibrium computed with

$\rho = 1$ are rewarded with additional gains that more than double the one achievable the ones relative to $\rho = 10$.

The right part of the Table 2 reports the largest difference between the equilibrium profits and the ones coming from a deviation. Such a worst-case measure ε^w , that quantifies the largest (over $50 \times 100 = 5000$ simulations) single increment in profits after departing from the equilibrium, is at least one order of magnitude larger than any ε . This clearly shows that there are instances in our 50 runs in which profits can be substantially increased by deviating from the strategy computed by ES in that specific experiment. This is, again, particularly true for $\rho = 1$, confirming that the accuracy of the equilibrium for this value of the recombination parameter is rather low.

Judging from the sizes of ε and ε^w , the quality of the equilibrium solutions are similar when $\rho = 5$ and $\rho = 10$, with a preference for the latter recombination parameter that often produces smaller additional profits.

Table 3: Fractions of cases, expressed as percentages, in which deviations are profitable and significant at various significance levels α .

		DEV>ES	$\alpha = 0.1$	$\alpha = 0.01$	$\alpha = 0.001$
$\rho = 1$	1	38.86	4.7	0.7	0.4
	2	25.36	2.0	0.2	<0.1
	3	7.58	0.7	0.1	<0.1
$\rho = 5$	1	30.00	3.0	0.4	0.2
	2	21.94	2.2	0.2	0.1
	3	15.11	1.0	0.1	<0.1
$\rho = 10$	1	24.68	2.4	0.3	0.2
	2	25.36	1.5	<0.1	<0.1
	3	6.59	0.5	<0.1	<0.1

Table 3 fine-tunes the information on the quality of the equilibrium evolved by our ES algorithm, displaying the fraction of cases in which deviations yield higher profits and how often the additional gain is statistically significant. Take again the row labelled “2” with $\rho = 5$: deviations are profitable in 21.94% of the simulations but the additional gains are significant at the 0.1, 0.01 and 0.001 levels only in 2.2%, 0.2% and 0.1% of the cases, respectively. Loosely speaking, with the exception of $\rho = 1$, that has already been shown to be a poor choice, most of the deviations do not increase the profits with respect to the optimal ES strategy and, even in the few cases in which mutations are profitable, the gains are only very rarely significant. There are, however, a small number of beneficial deviations that appears to produce significant gains, thus confirming the comments relative to the worst-case ε^w in Table 2.

5 Conclusion

We have analyzed in this paper the equilibrium strategies of heterogeneous agents in a CDA where limited time was available to trade. This form of time-related pressure induces agents to submit more aggressive orders as time goes by, according to a roughly linear schedule.

The equilibrium strategies were numerically computed decomposing the task to solve separate problems (one for every type of trader) and using an Evolution Strategies algorithm. Satisfactory results are robustly obtained when the recombination parameter is $\rho = 5$ and $\rho = 10$, but not when $\rho = 1$.

We evaluated the accuracy of our results using the concept of ε -equilibrium that allows to gauge to what extent strategic deviations are profitable. As it is impossible in our framework to exhaustively test all possible variations of the bidding functions, we resort to randomization sampling agents proportionally to their strength and selecting strategies to be mutated in three different ways. Even taking into account some potential underestimation, the incentive to deviate quantified by ε appears to be minute in all occasions, emphasizing the overall quality of the approximated equilibrium strategies calculated by Evolution Strategies.

References

- [Ashlock, 2006] Ashlock, D. (2006). *Evolutionary computation for modeling and optimization*. Springer.
- [Beyer and Schwefel, 2002] Beyer, H.-G. and Schwefel, H.-P. (2002). Evolution strategies: a comprehensive introduction. *Natural Computing*, 1:3–52.
- [Dawid, 1999] Dawid, H. (1999). On the convergence of genetic learning in a double auction market. *Journal of Economic Dynamics and Control*, 23(9-10):1545 – 1567.
- [Fano et al., 2011] Fano, S., LiCalzi, M., and Pellizzari, P. (2011). Convergence of outcomes and evolution of strategic behavior in double auctions. *Journal of Evolutionary Economics*, pages 1–26. 10.1007/s00191-011-0226-4.
- [Foucault, 1999] Foucault, T. (1999). Order flow composition and trading costs in a dynamic limit order market1. *Journal of Financial Markets*, 2(2):99–134.

- [Foucault et al., 2005] Foucault, T., Kadan, O., and Kandel, E. (2005). Limit order book as a market for liquidity. *Review of Financial Studies*, 18(4):1171–1217.
- [Parlour, 1998] Parlour, C. A. (1998). Price dynamics in limit order markets. *Review of Financial Studies*, 11(4):789–816.
- [Rosu, 2009] Rosu, I. (2009). A dynamic model of the limit order book. *Review of Financial Studies*, 22(11):4601–4641.
- [Rust et al., 1994] Rust, J., Miller, J. H., and Palmer, R. (1994). Characterizing effective trading strategies : Insights from a computerized double auction tournament. *Journal of Economic Dynamics and Control*, 18(1):61 – 96.