# An evolutionary multi-objective optimization of the market structure in call markets

Xinyang Li and Andreas Krause

***Abstract*— We evaluate an agent-based model featuring near-zero-intelligence traders operating in a call market with a wide range of trading rules governing the determination of prices, which orders are executed as well as a range of parameters regarding market intervention by market makers and the presence of informed traders. We optimize these trading rules using a multi-objective population-based incremental learning algorithm seeking to maximize the trading volume and minimize the bid-ask spread. Our results suggest that markets should choose a small tick size if concerns about the bid-ask spread are dominating and a large tick size if maximizing trading volume is the main aim. We also find that in contrast to trading rules in actual markets reverse time priority is an optimal priority rule.**

## I. Introduction

Market microstructure theory as used in conventional finance suggests that the trading rules applied by a market affect the prices at which trades occur, see [1] and [2] for an overview. This influence on prices should then also be visible in the statistical properties of returns such as their distribution and autocorrelations. In the highly structured models of market microstructure theory it is, however, difficult to evaluate a wide range of trading rules in a single model. Furthermore, the behavioral assumptions in those models make it difficult to assess the impact the changed trading rules have on the outcome, relative to behavioral influences.

In order to overcome these difficulties we develop an agent-based model in which traders use a very simple trading algorithm which does not assume rational behavior or any other optimizing rule. Such zero-intelligence (ZI) traders have been first introduced in [3] with the explicit aim to investigate the importance of the trading rules for the outcomes of trading. The strategic behavior has been considered to be a dominant influence factor for the market dynamics in previous research. However, [3] find that many of the major properties of double auction markets including the high allocative efficiency are primarily derived from the constraints imposed by the market mechanism, independent of traders' behavior. The zero intelligence approach has been widely applied in simulations of financial markets, particularly for analysis of the stylized facts, see [4], [5], [6], [7], [8], [9] and [10]. In [11] such traders have also been used to determine the optimal type of auction market. The use of appropriate automatons would allow us to focus on the influence the market structure, i.e. set of trading rules, has on the outcomes; in [10], models with near ZI traders do provide good return properties. In [12] and [13] a single-objective optimization of this model has been conducted using the trading-volume and bid-ask spread as objective functions and in this paper we extend this framework to a multi-objective setting to evaluate how any conflicts between different interests in market characteristics might be resolved.

With traders essentially behaving randomly with minimal restrictions, we are able to investigate a wide range of trading rules, e.g. the tick size, degree of intervention by market makers, priority rules, and market transparency, commonly found in financial markets and conduct research into the design of call markets to obtain the optimal combination of these trading rules.

Thus far only very limited attention has been paid to the optimization of financial markets in the agent-based literature. Apart from [12] and [13], most notably [11] and [14] investigate optimization of the real-valued parameter set for adaptation in the trading agents using Genetic Algorithms (GA), an adaptive heuristic search algorithm with the evolutionary ideas of natural selection and genetic recombination. It is shown that a good result is reported with random initial value using GA together with an appropriate function. In addition, a recently developed framework of optimization, the population-based incremental learning (PBIL) has been widely applied. It is a type of evolutionary algorithm in which the genotype of the whole population is evolved rather than individual chromosomes. This algorithm, proposed by [15], has been found to be simpler and to achieve better results than the standard genetic algorithm in many circumstances, e.g. [16]. Therefore, in this paper, we apply the PBIL as the optimization algorithm and extend its use to a multi-objective setting by using the trading volume and bid-ask spread as objective functions.

Using the results obtained from this research it is possible to derive recommendations to exchanges, regulators on establishing the optimal market structure, for securities issuers to choose the best exchange for their listing and for investors to choose the most suitable exchange for trading.

We continue in section II by introducing our model as well as the trading rules considered. Section III then discusses the results of our computer experiments and section IV concludes the findings.

## II. Description of the market

### A. The behavior of traders

We investigate a market in which a fixed number of $N$ traders trade a single asset in a call market. At any time each

Xinyang Li is with the School of Management, University of Bath, Bath BA2 7AY, Great Britain, E-Mail: xl212@bath.ac.uk

Andreas Krause is with the School of Management, University of Bath, Bath BA2 7AY, Great Britain, E-Mail: mnsak@bath.ac.uk

trader is either a buyer or seller of the asset and submits buy orders $B_i$, $i = 1, \ldots, N$ such that at time $t$ the limit price is taken from a log-normal distribution:

$$\ln B_i^t \sim iidN\left(\ln \bar{P}_t + \mu_{buy}, \sigma_{buy}^2\right), \tag{1}$$

where $\bar{P}_t$ is the long-term fundamental value in time period $t$, which we here assume to be equal to the initial price $P_0$.[1] $\mu_{buy}$ denotes the average amount by which the bid price exceeds the fundamental value, and $\sigma_{buy}^2$ represents the variance of bid prices around the mean. With $\mu_{buy} < 0$ the limit bid price will on average be below the fundamental value, although traders may well submit orders with limit prices above the fundamental value given the random nature of the limit price in our model. We might interpret this either as uncertainty about the fundamental value to which traders pay limited attention, different opinions about the true fundamental value or the fact that many traders will ignore the fundamental value to a large degree in their decision-making. While experiments have shown that the exact specification of the decision-making process is not affecting results, we require a minimal amount of information which traders use as a common anchor for their decision; this is necessary to avoid the limit prices and thereby transaction prices to evolve such that an infinitely large bubble emerges. This constraint on the behavior of traders thus implicitly acts as a budget constraint as too large limit prices are not permitted and similarly too small limit prices will not be observed, acting as a minimum size requirement for entering the market.

If we denote by $\widehat{P}_i^{t-1}$ the price at which a trader bought the asset the last time, the limit price of a sell order is chosen according to

$$\ln S_i^t \sim iidN\left(\ln \widehat{P}_i^{t-1} + \mu_{sell}, \sigma_{sell}^2\right), \tag{2}$$

in which $\mu_{sell}$ denotes the average amount by which the ask price exceeds the price previously paid by the trader, and $\sigma_{sell}^2$ represents the variance of ask prices. A trader will only be able to sell those shares he actually holds, i.e. we do not allow for any short sales, thereby acting implicitly as a budget constraint on the behavior of traders.

The order size for a sell order will always be equal to the number of shares held. The order size for buy orders $Q_i^t$ is a random variable with

$$\ln Q_i^t \sim iidN\left(\mu_{size}, \sigma_{size}^2\right), \tag{3}$$

where $\mu_{size}$ denotes the average of the order size, and $\sigma_{size}^2$ is the variance of the order size.

An order remains in the order book until it is filled or canceled; for partially filled orders the remainder of the order remains in the order book. An order not filled after $T_i^t$ time steps is canceled, where

$$\ln T_i^t \sim iidN\left(\tau, \sigma_\tau^2\right), \tag{4}$$

[1]We could also introduce a positive long-term trend of the fundamental value without changing the results of our model.

in which $\tau$ is the average time of order remains in the order book, and $\sigma_\tau^2$ denotes the variance of this time.

The canceled order is replaced by a new order taken from the following distributions:

$$\begin{aligned} \ln B_i^t &\sim iidN\left(\ln \bar{P}_t + \mu_{buy}, \sigma_{buy}^2\right), \\ \ln S_i^t &\sim iidN\left(\ln P_i^{t-1} + \mu_{sell}, \sigma_{sell}^2\right), \end{aligned} \tag{5}$$

where $P_i^t$ denotes the market price at time $t$.

Whether a trader is a buyer or a seller is determined as follows: if his last transaction was to buy the asset he becomes a seller and if his last transaction was to sell the asset he becomes a buyer. A change from buyer to seller or vice versa only occurs if he has no order remaining in the order book. In the initialization of the experiments buyers and sellers are determined randomly.

### B. Determination of transaction prices

Following the price formation approach applied in [3] and [11], the transaction price is determined where the demand and supply curves intersect, i.e., the price at the maximal trading volume is chosen as the trading price. In this market, limit orders with the highest bid prices are first traded and cleared in the market; oppositely, the cheapest sell orders are traded with priority. If we find that there are multiple prices at which the trading volume shows the same maximal value, we employ trading rules to determine which of the prices will be chosen. Any imbalances between buy and sell orders at the transaction price will lead to the need for rationing; how this rationing of buy or sell orders is conducted will depend on the trading rules as outlined below.

In contrast to the models in [3] and [11], however, we do not use a continuous double auction market but rather a call market in which orders of traders batched. We batch all orders in each time step, where a time step consists of the submission and revision of orders as well as the batching of orders, the determination of the transaction price and execution of the trades.

### C. Trading rules considered

*1) Tick size:* In the market we are able to vary a wide range of trading rules. We will firstly investigate different *tick sizes*, i.e. minimum differences between prices at which orders can be submitted. The tick size has several impacts during trading. As it represents the cost of getting inside other competitors quote, the tick size affects the motivation of submitting limit orders. In addition, the tick size has an impact on the spreads. As reported in many empirical investigations, e.g. [17] and [18], amongst others, the bid-ask spread declined dramatically by about 25% with the reduction of tick size from 1/8 to 1/16 dollar on the New York Stock Exchange (NYSE). In order to make limit prices to comply to the tick size, we will lower any limit price of buy orders as determined in (1) and (5) to the next permissible price and similarly raise the limit price of sell orders determined by (2) and (5) to the next permissible price.

*2) Priority rules:* Secondly, different *priority rules* are employed to determine the rationing of orders in the case of an imbalance between buy and sell orders at the transaction price, see [19] and [20] for an overview of the different priority rules found in several markets. The enforcement of priority rules, as the primary difference between market structures, is another important design feature of trading systems. We use in particular *time priority*, which is the most commonly used rule. It adheres to the principle of first-come first-served, and ensures that orders submitted earlier will be filled first; *reverse time priority* in which orders submitted later will receive priority to be filled; another frequently used rule to promote traders to place larger orders is the *size priority* in which larger orders receive priority; *random selection* in which the orders to be filled are selected randomly and with *pro-rata selection*, a common practice on many financial market such as the Stock Exchange of Hong Kong, the old Toronto Stock Exchange and the batch systems, in which all orders get filled partially to the same fraction.

*3) Multiple prices:* Thirdly, for the case of *multiple prices* at which the trading volume is maximal we determine the transaction price to be either the price closest to the previous price, the price furthest from the previous price, the highest price, the lowest price, the price with minimum order imbalance (the difference between the volume of buy and sell orders at the transaction price), the price with maximum order imbalance or a randomly selected price.

*4) Market transparency:* Fourthly, we also consider *market transparency*, which is defined by [1] as "the ability of market participants to observe the information in the trading process". In this context, information refers to knowledge about the prices, the size and direction of orders, and the identities of market participants. In a transparent market, traders are able to have access to information on the order book and react to any orders submitted by other traders. This could reduce the magnitude of adverse-selection problems. Hence transparency is expected to increase profit for traders. [21] show that transparency does reduce the trading costs incurred by uninformed traders in theory. Empirically, [22] discover that the pre-trade transparency can narrow the spreads. However, one can argue that transparency can make it difficult to supply liquidity to large traders, who may be reluctant to submit limit orders, since the disclosure may convey information which makes the price moves against the trader′s position.

In order to replicate this aspect of the market we assume that a fraction of $\gamma$ of the traders has access to the order book and can observe the potential transaction price as well as the ensuing order imbalance if the trades were to happen instantly. They use this information to revise their own order size according to the size of the order imbalance $\delta$ for a buy and sell order, respectively:

$$\begin{aligned} \widehat{Q}_i^t &= Q_i^t - \alpha\delta, \\ \widehat{Q}_i^t &= Q_i^t + \alpha\delta, \end{aligned} \tag{6}$$

where $\alpha$ represents the fraction of order size revised, $Q_i^t$ is the order size before revision, and $\widehat{Q}_i^t$ is the order size after revision. This revised size is then used to determine the transaction price.

*5) Market making:* As a final aspect we consider the *intervention of a market maker* into the trading process. A market maker would intervene or influence the prices such that he is prepared to trade a fraction $\theta$ of the order imbalance at any time in the market with the existence of imbalance between demand and supply at the transaction price by submitting an offsetting order with price

$$\widehat{\widehat{P}}^t = P^t + \lambda I^t, \tag{7}$$

where $I^t$ denotes the inventory of the market maker, i.e. the number of shares held by him, $\lambda$ is the price adjustment of market maker. Holding a volume of shares as inventory, the $I^t$ is a positive number; on the other hand, the $I^t$ is a negative value representing a purchasing position. Therefore, the bid price set by the market maker would be lower than the market price, while the ask price is higher than the market price. Such a linear relationship between the price and inventory has been established in the inventory-based models of market-making, see e.g. [23], [24] or [25], although other mechanisms have been proposed in the literature, e.g. in [26].

### D. Optimization of market structures

The methodology used to optimize the market structure is a computer experiment in which trading is simulated over a given number of time periods with a given market structure. The optimization of the trading rules is conducted evolutionary by population-based incremental learning (PBIL), using the trading volume (maximizing) and the bid-ask spread (minimizing) as our performance functions for our multi-objective optimization. The PBIL, described by [15], is "a method of combining genetic algorithms and competitive learning for function optimization."

*1) Genetic algorithm:* Genetic algorithms (GA), developed in [27], are automated optimization algorithms based upon the principles of natural selection and genetic recombination. GAs maintain a group of potential solutions to the objective function being optimized, a so called "population", constructed randomly for the initial generation. Each population member is referred to be a chromosome, represented by a string of binary alphabet. In each generation, the fitness of each chromosome or potential solution is estimated, i.e. it measures how well each solution optimizes the objective function. The better fitted sets are selected as "parents" and randomly paired to create a new optimal set as the next generation with the reproduction process of crossover and mutation. Since parents with higher fitness are more likely to pass their characteristics on to the child, the average fitness would increase. Therefore, the potential solutions in the last generation are expected to be the best solutions found for the optimization.

Although GAs have been applied successfully in financial market research, see [11] and [14], many issues, such as efficient problem representation and adequate scaling of functions to make sure the good genes pass down to the next generation, need to be solved. Among a fixed number of generations, it is very difficult for the GA to return the optimal solution due to the randomized searching, and the inability of comparing the often very small difference between good and optimal solutions.

*2) Competitive learning:* Competitive learning, clustering a number of unlabeled points into different groups based on the similarity of points, is often used in the field of artificial neural networks. The input cells in this network represent the feature vector for each point; the outputs represent the group where the point has been allocated into. With random initial weights $w_{ij}$, the activation of the output neuron $i$ is computed during learning with the following formula

$$output_i = \sum_j w_{ij} * input_j \tag{8}$$

where $w_{ij}$ is the weight of the connection between input unit $i$ and output unit $j$. The weights of the output unit with the highest activation, i.e. the winning neuron, are modified to represent the current input better:

$$\Delta w_{ij} = \eta(input_j - w_{ij}) \tag{9}$$

with $\eta$ being the learning rate. The learning process would repeat until the network has stabilized. In this framework, all input units get the same signal, while only the winning neuron with the highest output signal is allowed to fire and used to represent the current input. After the learning process is complete, the weights for each output neuron can be viewed as a prototype vector for a group represented by this neuron.

*3) Population-based incremental learning:* In analogy to GAs, the population-based incremental learning(PBIL) algorithm maintains a population of potential solutions evolving over a number of generations. However, the PBIL attempts to create a probability vector, measuring the probability of each bit position having a ”1” in a binary solution string, to define a population of a genetic algorithm. Instead of transforming each individual into a probability vector used for generating and recombination, the probability vector is moved towards representing the high evaluation vectors with a similar manner to the competitive learning process. The probability vector $\pi_t$ is updated based on the following rule

$$\pi_t = (1-\eta)\pi^*_{t-1} + \eta\widehat{v} \tag{10}$$

where $\pi^*_{t-1}$ denotes the probability of containing a 1 in each bit position that was used in the previous time period, and $\widehat{v}$ represents the best solution in the current generation selected according to the fitness function of the optimization and $\eta$ the learning rate. The probabilities are subject to mutation at a mutation rate $\xi$ and the actually chosen probability $\pi^*_t$ will be

$$\pi^*_t = (1-\xi)\pi_t + \xi\varepsilon \tag{11}$$

with $\varepsilon \sim U[-1;1]$ and $\pi^*_t$ restricted between 1 and 0. This algorithm is capable to maintain diversity in search as the same probability vector could generate distinct population.

In [15], the author compares empirical performance of a standard genetic algorithm to the simple PBIL, and shows that the PBIL is capable to attain the results more accurately and faster than a standard GA. As a combination of genetic algorithms and competitive learning it makes PBIL an successful and efficient search mechanism employed in such complex optimization problems as the one presented in this paper.

### *E. Multi-objective optimization*

Applications of multi-objective optimization using the PBIL algorithm are rare thus far, e.g. [28], and this paper is one of the few applications in this field. In multi-objective optimization no single best solution typically exists but in most cases a trade-off between the different objective functions are observed. For this reason the reference vector $\widehat{v}$ in equation (10) will be determined as the Pareto-efficient solution closest to the currently used vector. ”Closest” is defined by the Euclidean norm and a ”Pareto-efficient” solution is one for which there is no solution which is superior for both objective functions.

To employ the PBIL algorithm, we first create a probability vector specifying the probability of each bit position having the value one. All the initial value for the PBIL process is determined randomly, i.e. the initial probability of each bit position is considered as 0.5. With the probability vector, a population could be produced accordingly. In each generation we determine the fitness of each potential solution and the best fitted set in the current generation is selected and used to update the probability vector for the subsequent generation using (10) and (11).

In each time step we determine 100 different parameter constellations using $\pi^*_t$ and then determine the best performing parameter constellation from these 100 different market simulations that then makes $\widehat{v}$.

Each trading rule is coded into a vector $v$, where the precision of the continuous variables $\alpha$, $\lambda$, $\gamma$, $\theta$ is such that each variable is divided into 17 bits each, the tick size $t$ into 20 bits; the discrete variables (priority rules, multiple prices) are coded such that all rules are covered.

As is common with multi-objective optimization, we do not observe an easy convergence of results (even after 5,000 generations no convergence towards a clearly identifiable Pareto-efficient frontier was observed). For this reason we run the optimization for 500 generations and use the entire population of the resulting final generation to analyze our results. This length is about 4 times the length it took the single-objective optimization in [12] and [13] to converge and should therefore represent an adequate time length for the evolutionary algorithm to evolve.

## III. Results of computer experiments

### *A. Parameter constellations considered*

We consider a market with 100 traders, which consist of 50 buyers and 50 sellers for the first round. The order book contains the traders' ID number, whether they are buying or

TABLE I

PARAMETER VALUES CONSIDERED IN THE COMPUTER EXPERIMENT

| Parameter | Description | Value |
|---|---|---|
| $\mu_{buy}$ | Mean of bid exceeding fundamental value | -0.1 |
| $\mu_{sell}$ | Mean of ask exceeding fundamental value | 0.1 |
| $\sigma_{buy}$ | Standard deviation of bid | 0.1 |
| $\sigma_{sell}$ | Standard deviation of ask | 0.1 |
| $\tau$ | Mean time of order in order book | 1+ln100 |
| $\sigma_{\tau}$ | Standard deviation of time in order book | 1 |
| $\mu_{size}$ | Mean of order sizes | 1 |
| $\sigma_{size}$ | Standard deviation of order sizes | 1 |
| $\mu$ | Trend of the fundamental value | 0 |
| $t$ | Tick size | [0,10]* |
| $\alpha$ | Fraction of order size revised | [0,1]* |
| $\lambda$ | Price adjustment of market maker | [0,1]* |
| $\gamma$ | Fraction of informed traders | [0,1]* |
| $\theta$ | Imbalance traded by market maker | [0,1]* |

* denotes the parameter range used in the optimization.

selling, their limit price, order size, order submission time and length until the order is to be revised. The initial order book is constructed randomly using the parameter settings described in table I and the initial price $P_0$ set at 100. We assume that the trading price equals the previous price if there is no trading.

Each simulation is run for 2,000 time steps, where the first 1,000 data is eliminated for the investigation. The multi-objective PBIL optimization is conducted using a population size of 100 over 500 generations with learning rate of 0.2 and a mutation rate of 0.01. We repeat the multi-objective optimization 20 times with the same parameter constellation to reduce the amount of noise remaining from the lack of convergence.

### B. Evaluation of computer experiments

In figure 1 we show the trading volume and spread of the final generation for the entire population for all 20 runs of our computer experiments restricted to the area close to the Pareto-efficient frontier. We observe a trade-off between the maximal trading volume and the minimal spread, the approximate location of the Pareto-efficient frontier is sketched by the line to the lower right. This figure clearly shows that a low spread will be associated with a low trading volume while a high trading volume will necessitate a large spread. The optimal combination between trading volume and spread will depend on the preferences of the decision-maker and how he values the importance of these two aspects.

From figure 1 we have identified 8 markets that approximately determine the Pareto-efficient frontier, these points are identified as large red points and associated with numbers; the market structures of these eight markets are shown in figure 2. The first point to notice is that unless the decision-maker puts great emphasis on a large trading volume at the expense of the spread, the tick size chosen should be very small as can easily be observed from the top left panel. Results are more ambiguous for the fraction of informed traders (top right panel), here no clear pattern can be identified for different points along the Pareto-efficient frontier, suggesting that results are not very sensitive to this variable and informed traders should ideally make approximately 5% of the trader population. Similarly the fraction of orders revised by those informed traders seem to be independent of the location on the Pareto-efficient frontier at approximately 0.6. The same result holds for the intervention of the market maker who should trade approximately 60% of the order imbalance and a price adjustment fluctuating widely with an average of about 0.3. Apart from one instant associated with a high trading volume and large spread, the optimal priority rule is reverse time priority and the optimal multi-price rule is to choose the nearest price to the previous transaction price. At this stage we can summarize our results in stating that for determining the location on the Pareto-efficient frontier only the tick size seems to have a significant impact.

Comparing these results with those in [12] and [13] who use the same model and optimization technique, although applying only a single objective function - the trading volume and spread, respectively - we observe that our results are largely consistent with those obtained in these papers. We find that most variables take values similar to the optimal values in the single-objective optimization, where many were found to be very variable when running different experiments. Furthermore we are able to recover one feature on the tick size that has also been found in those papers, when maximizing trading volume it was found that a large tick size was optimal while when minimizing the spread a small tick size was preferable. This feature we also found in our results as we have established above that when concerns about the spread dominate the tick size should be small, while it should be large when the trading volume is of more importance.

Another result was also found to be stable in multi-objective optimization: the optimal priority rule is reverse time priority, i.e. orders submitted later are executed earlier. This result was established when minimizing the spread in [13] and is confirmed here in a multi-objective setting; when maximizing the trading volume in [12] priority rules were found to be ambiguous in the optimal market structure, a fact that might cause the time priority to be optimal in one case. Reverse time priority is in opposition to most real markets that apply price priority and should be an aspect that might be beneficial for market to consider if they seek to optimize their market structures.

## IV. CONCLUSIONS

In this paper we investigate the combination of a wide range of trading rules in a multi-objective optimization by employing population-based incremental learning (PBIL) in call markets seeking to maximize trading volume and minimize the bid-ask spread. As trading rules we include the tick size, priority rules, multi-price rule, intervention of market makers and market transparency. In order to eliminate the influence of complex trader behavior we use an agent-based model in which traders behave nearly randomly, such that any properties arising can be attributed to the impact of the trading rules directly rather than trader behavior.

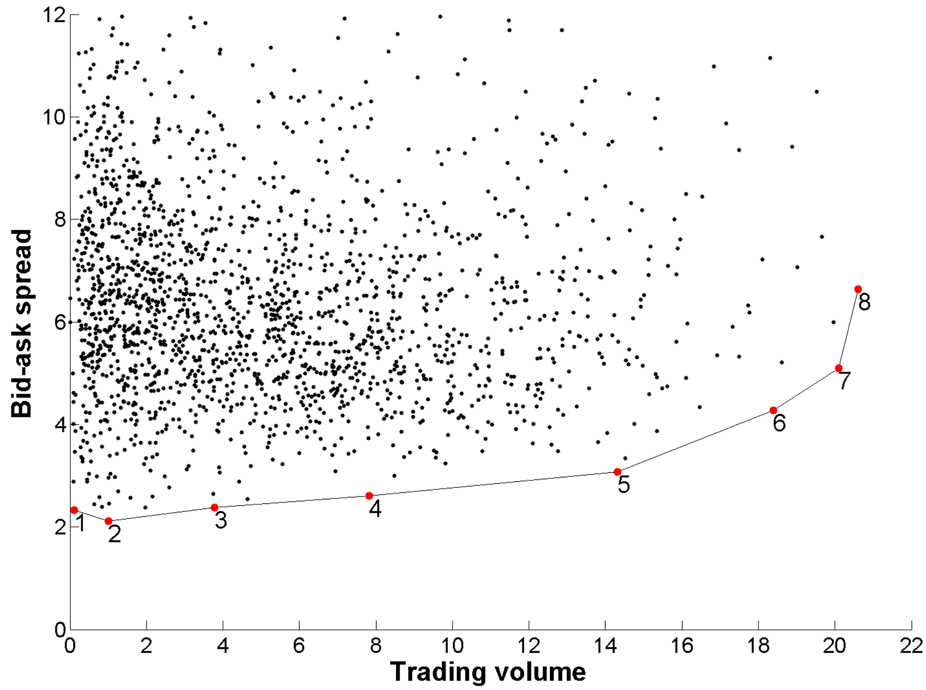


Fig. 1. Market performance after 500 generations

Conducting such an analysis we analyze the market structures of those markets close to the Pareto-efficient frontier. The results show that when concerns about the bid-ask spread are dominating a small tick size should be chosen and a large tick size if concerns about trading volume are more important. We also find that as a priority rule markets should use reverse time priority rather than time priority as currently done in nearly all markets. These results have direct consequences for the optimal design of financial markets in terms of maximization of trading volume while ensuring a low bid-ask spread, and thus might inform any market reforms considered by stock, bond or derivatives markets.

In future research the proposed framework can easily be extended to include other objective functions, like minimizing volatility, maximizing share value, or maximizing the trading profits to small traders as alternative or additional objective functions. Such research would allow us to balance a wider range of interest in the market and investigate the sensitivity of the optimal trading rules to the different preferences of decision-makers, thereby giving a more complete picture of the influences on market performance.

## References

[1] Maureen O'Hara. *Market Microstructure Theory*. Blackwell, Oxford, Uk, 1995.

[2] Ananth Madhavan. Market microstructure: A survey. *Journal of Financial Markets*, 3:205–258, 2000.

[3] Dhananjay K. Gode and Shyam Sunder. Allocative efficiency of markets with zero-intelligence traders: market as a partial substitute for individual rationality. *Journal of Political Economy*, 101:119–137, 1993.

[4] Paczuski M. Shubik M. Bak, P. Price variations in a stock market with many agents. *Physica A*, 246:430–453, 1997.

[5] Sergei Maslov. Simple model of a limit order-driven market. *Physica A*, 278:571–578, 2000.

[6] Damien Challet and Robin Stinchcombe. Analyzing and modeling /1+1d markets. *Physica A*, 300:285–299, 2001.

[7] Andreas Krause. Fat tails and multi-scaling in a simple model of limit order markets. *Phisica A*, 368:183–190, 2006.

[8] Dan Ladley and Klaus Reiner Schenk-Hoppe. Do stylised facts of order book markets need strategic behaviour? Swiss finance institute research paper series 07-20, Swiss Finance Institute, 2007.

[9] Abraham Othman. Zero-intelligence agents in prediction markets. In *Proceedings of the Seventh International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, pages 879–886. IFAAMAS, 2008.

[10] Xinyang Li and Andreas Krause. *IDEAL 2009, Lecture Notes in Compter Science 5788*, chapter A comparison of market structure with Near-Zero-Interlligence traders, pages 703–710. Springer-Verlag, 2009.

[11] Dave Cliff and Janet Bruten. Minimal intelligence agents for bargaining behaviors in market-based environments. HP Lab Report HPL-97-91, HP, 2001.

[12] Xinyang Li and Andreas Krause. Determining the optimal market structure using near-zero intelligence traders. Technical report, University of Bath, 2009.

[13] Xinyang Li and Andreas Krause. Determining the optimal market structure using near-zero intelligence traders. Technical report, University of Bath, 2009.

[14] Dave Cliff. Explorations in evolutionary design of online auction market mechanisms. *Electronic Commerce Research and Applications*, 2:162–175, 2003.

[15] Shumeet Baluja. Population-based incremental learning: A method for integrating genetic search based function optimization and competitive learning. Technical report cmu-cs-94-163, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, 1994.

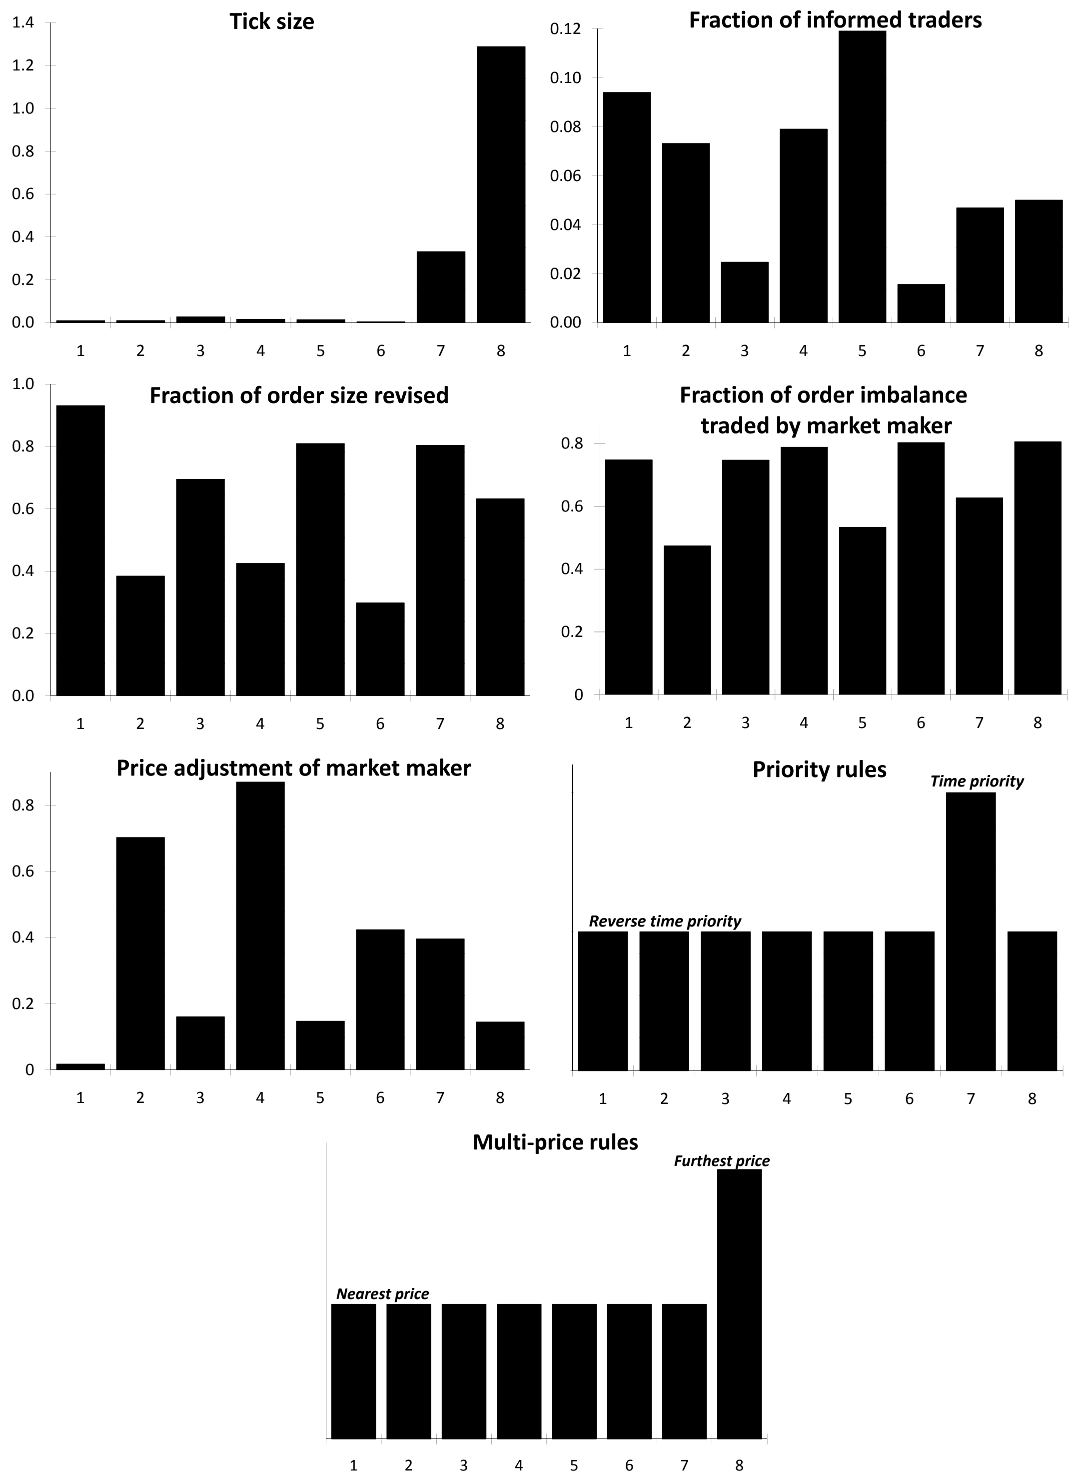


Fig. 2. Market structures near the Pareto-efficient frontier


[16] Shummet Baluja. An empirical comparison of seven iterative and evolutionary function optimization heuristics. Technical reports, Carnegie Mellon University, 1995. Available via. anonymous ftp at: reports.adm.cs.cmu.edu.CMU-CS-95-193.

[17] Michael A. Goldstein and Kenneth A. Kavajecz. Eighths, sixteenths, and market depth: Changes in tick size and liquidity provision on the nyse. *Journal of Financial Economics*, 56:125–149, 2000.

[18] Charles M. Jones and Marc L. Lipson. Sixteenths: direct evidence on institutional execution costs. *Journal of Financial Economics*, 59(2):253–278, February 2001.

[19] Robert Alan Schwartz. *Equity markets: Structure, trading, and performance*. New York: Harper & Row, 1988.

[20] I. Domowitz. A taxonomy of automated trade execution systems. *Jounal of International Money and Finance*, 12:607–631, 1993.

[21] M. Pagano and A. Rell. Transparency and liquidity: A comparison of auction and dealer markets with informed trading. *The Journal of*

*Finance*, 1996.

[22] C.G. Koedijk M.D. Flood, R. Huisman and R.J. Mahieu. Quote disclosure and price discovery in multiple dealer financial markets. *Review of Financial Studies*, 12:37–59, 1999.

[23] H. R. Stoll. The supply of dealer services in securities markets. *Journal of Finance*, 33:1133–1151, 1978.

[24] T. Y. Ho and H. R. Stoll. On dealer markets under competition. *Journal of Finance*, 35:259–267, 1980.

[25] T. Y. Ho and H. R. Stoll. The dynamics of dealer markets under competition. *Journal of Finance*, 38:1053–1074, 1983.

[26] Nils H. Hakansson, Avraham Beja, and Jivendra Kale. On the feasibility of automated market making by a programmed specialist. *Journal of Finance*, 40:1–20, 1985.

[27] John Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, 1975.

[28] Sujin Bureerat and Krit Sriworamas. *Soft computing in industrial applications*, volume 39/2007, chapter Population-based incremental learning for multiobjective optimisation, pages 223–232. Springer, 2007.

[29] Giovanni Dosi Giulio Bottazzi and Igor Rebesco. Institutional architectures and behavioral ecologies in the dynamics of financial markets. *Journal of Mathematical Economics*, 41:197–228, 2005.

[30] Rama Cont. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1:223–236, 2001.

[31] Paolo Pellizzari and Arianna Dal Forno. Comparison of different trading protocols in an agent-based market. *Journal of Economic Interaction and Coordination*, 2:27–43, 2007.