Evolutionary Optimization of Parameter Sets for Adaptive Software-Agent Traders in Continuous Double Auction Markets

Dave Cliff
Publishing Systems and Solutions Laboratory
HP Laboratories Bristol
HPL-2001-99
April 26th , 2001*

E-mail: cliffd@hplb.hpl.hp.com

auctions, agents, traders, zip, genetic algorithms

The continuous double-auction (CDA) is a powerful market mechanism, noted for its speed and efficiency, and is the mechanism underlying the organization of open-outcry 'trading pits' at major international derivatives markets. In previous publications, Cliff & Bruten demonstrated that software 'trading agents' need more than zero intelligence to give human-like CDA price-equilibration behavior and presented results from experiments with simple adaptive trading agents in CDA markets and in one-sided auctions. These agents give very good performance on standard measures of trading activity such as allocative efficiency, Smith's 'alpha' measure of price convergence, and profit dispersion, but only when parameters governing the adaptation mechanism are set to appropriate values. Determining good or optimal combinations of parameters by hand is possible, but can be labor-intensive. This paper presents the first results from using a genetic algorithm to optimize all the real-valued parameters governing adaptation in the trading agents. It is shown that a simple genetic algorithm (GA), in combination with an appropriate evaluation function, can deliver good parameter settings from random initial-value conditions. The evolutionary trajectories of the population through the 8-dimensional parameter space are illustrated, and the use of the GA to identify parameters that are redundant or even harmful is discussed.

* Internal Accession Date Only Approved for External Publication
Presented at Workshop on Computional Markets, Agents '98 Conference, Minneapolis, May 1998

Evolutionary Optimization of Parameter Sets for Adaptive Software-Agent Traders in Continuous Double Auction Markets

Dave Cliff*
Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge MA 02139, U.S.A.
davec@ai.mit.edu

Abstract

The continuous double-auction (CDA) is a powerful market mechanism, noted for its speed and efficiency [22], and is the mechanism underlying the organization of open-outcry 'trading pits' at major international derivatives markets. In previous publications, Cliff & Bruten demonstrated that software 'trading agents' need more than zero intelligence to give human-like CDA price-equilibration behavior [10, 13]. Cliff and Bruten presented results from experiments with simple adaptive trading agents in CDA markets [10, 12] and in one-sided "retail market" auctions [15]. These agents give very good performance on standard measures of trading activity such as allocative efficiency, Smith's [42] 'α' measure of price convergence, and profit dispersion, but only when parameters governing the adaptation mechanism are set to appropriate values. Determining good or optimal combinations of parameters by hand is possible, but can be labor-intensive. This paper[1] presents the first results from using a genetic algorithm to optimize all the real-valued parameters governing adaptation in the trading agents. It is shown that a simple genetic algorithm (GA), in combination with an appropriate evaluation function, can deliver good parameter settings from random initial-value conditions. The evolutionary trajectories of the population through the 8-dimensional parameter space are illustrated, and the use of the GA to identify parameters that are redundant (or even 'harmful') is discussed.

1 Background

The classical theoretical picture of equilibration (price formation at a competitive equilibrium) dictates that the number of trading agents (buyers and sellers) in the market is practically infinite, or very large at least. Yet, in a series of experiments commencing in the late 1950's, Smith (e.g. [42]) demonstrated that markets consisting of surprisingly small numbers of human traders could rapidly converge on the theoretical equilibrium price given by the intersection of the market's supply and demand curves. Smith's experimental results have been widely replicated and extended, and it is now generally accepted that stable equilibria can be reached with fewer than twenty traders.

*Current address: Digital Media Systems Dept, PSSL, Hewlett-Packard Labs, Bristol BS34 8QZ, UK. cliffd@hplb.hpl.hp.com.

[1] This paper is revised and significantly extended from [11].

In Smith's experiments, human subjects were told to be either buyers or sellers. Each seller was given a number of units of an arbitrary commodity to sell and each buyer was given the right to buy some units, and some currency. For each unit, each trader was given a *limit price* that was private (known only to that trader). A buyer couldn't pay more than her limit price for a unit of the commodity, and a seller couldn't sell a unit for less than her limit price. Typically, different traders had different limit prices: the distribution of limit prices determined the market supply and demand curves. In each experiment, time was divided into discrete periods known as 'days': most experiments consisted of 5 to 10 consecutive 'days'. At the start of each day, the rights to buy or sell units of commodity were distributed between the subjects. Each day ended either when no more traders were willing or able to trade, or when a pre-set time-limit expired (typically a 'day' lasted 5–10 minutes). During each day, the traders operated within a specific market structure: in many of Smith's experiments, the CDA was used, but he also experimented with one-sided auction 'retail' markets (where only sellers quote prices). In the early experiments, the traders operated in the experimental equivalent of an open-outcry 'trading pit', but Smith subsequently developed methods where the traders communicated with each other via a network of computer terminals. Smith's work helped establish the field now known as *Experimental Economics* [43, 19, 31]. The intention of the work described here is to develop artificially-intelligent software trading-agents with the bargaining capabilities necessary for small groups of traders to show price-equilibration in market scenarios similar to those studied by Smith.

Although it may seem intuitively obvious that some form of 'intelligence' or adaptation is necessary in trader-agents, Gode & Sunder [23] presented results that appear to indicate that *zero-intelligence* agents can exhibit human-like behavior in CDA markets. Gode & Sunder's zero-intelligence trading agents simply generated random prices for bids or offers, subject to the constraint that they could not enter into loss-making deals. Gode & Sunder's work has proven influential.[2] However, Cliff & Bruten [10, 13] demonstrated that Gode & Sunder's result only holds in very specific circumstances and that, in general, some 'intelligence' in the form of adaptivity or sensitivity to previous and current events in the market is necessary. Consequently, Cliff & Bruten [10, 12] described simple trading agents with adaptive capabilities. They refer to these as "zero-intelligence-plus" (ZIP) traders. The ZIP traders' basic adaptive mechanisms are elementary machine-learning techniques.

The emphasis in Cliff & Bruten's work was on creating *simple* autonomous software agents for bargaining in market-based environments. This emphasis on simplicity came not only from a desire for computational efficiency (important if hundreds or thousands of such agents are active on a network), but also from a desire to speculatively sketch the minimum mechanistic complexity necessary and sufficient for explaining human bargaining behaviors in specific market environments. There are at least three significant potential applications of bargaining agents such as ZIP traders: internet-based trading and e-commerce; economic modeling; and market-based control (e.g. [9]) — these issues are discussed further by Cliff [10], who also gives a complete description of the design of ZIP trading agents, shows results from many experiments in different styles of market environment, and includes all the C source-code for the system. Two recent theses have further explored the use of ZIP traders. The first, by van Montfort [48], discusses experiences in using these ZIP traders in markets where there may be potentially hundreds or thousands of agents, where there

[2] For example, the following 30 texts approvingly cite Gode & Sunder's 1993 *JPE* paper on ZI traders: [3, pp.230–231], [5, pp.253, 258], [21, p.19], [22, p.*xxiii*], [32, pp.292, 294], [40, pp.160–161, 175], [19, p.132], [33, p.2], [26, p.1082], [44, p.310], [1, p.186], [4, p.674] [27, p.228], [28, p.370], [30, pp.570, 580], [34, p.226], [38, pp.52–55, 80–81], [46, p.475], [6, pp.1318, 1333], [17, p.678], [18, p.383], [29, p.276], [36, p.266], [41, p.32], [45, p.2], [47, p.461], [7, p.320], [8, pp.183–184], [20, p.623], and [24, pp.604–605].

may be spatial structure or segmentation in the market (e.g., the traders are distributed over some space, and each trader can only transact with other traders in its local neighborhood), and where agents are permitted to engage in arbitrage. In the second thesis, by van Tol [49], the ZIP-trader specification was extended to a generalized form capable of dealing with variations in the amount of market information (i.e., the number of quote-prices) employed by the trader's pricing strategy.

The present paper describes the use of a genetic algorithm to optimize the values of real-valued numeric parameters governing adaptation or learning in ZIP traders. Before these genetic optimization experiments can be described, it is necessary to introduce the ZIP trading strategies and adaptation mechanisms in more detail.

2 ZIP Traders

Each ZIP trader operates by maintaining a *profit margin* that it uses for calculating the price it 'quotes' (offers or bids) in the market: the profit margin determines the difference between the price the agent quotes and that agent's limit price for the unit of commodity the agent is trading. The 'aim' of each ZIP agent is to maximize profit generated by trading in the market. Thus, the problem of designing a trading agent can be considered as a combination of two issues: the *qualitative* issue of deciding *when* to increase or decrease the profit margin, and the *quantitative* issue of deciding *by how much* the margin should be altered.

For reasons discussed in detail by Cliff [10], each ZIP trader makes the qualitative decision of when to alter its margin on the basis of four factors. The first factor is whether the agent is *active* in the market: agents are active until they have sold or bought their full entitlement of units of the commodity. The remaining three factors are data concerning the last quote by any agent in the market, which we refer to as Q. Each ZIP trader notes whether Q was an offer or a bid, whether Q was accepted (i.e., led to a transaction) or rejected (ignored by the traders in the market), and whether Q's price, $q(t)$, is greater than or less than the price the ZIP trader would currently quote. The price a ZIP trader i would quote at time t is referred to as that trader's *quote-price*, $p_i(t)$, which is calculated from i's limit price $\lambda_{i,j}$ (for i's jth unit of commodity) and i's current profit coefficient $\mu_i(t)$ using $p_i(t) = \lambda_{i,j}(1 + \mu_i(t))$. This implies that a seller's margin is raised by increasing μ_i and lowered by decreasing μ_i, with the constraint that $\mu_i(t) \in [0, \infty); \forall t$. The situation is reversed for buyers: they raise their margin by decreasing μ_i and lower it by increasing μ_i, subject to $\mu_i(t) \in [-1, 0]; \forall t$. Note that if trader j was responsible for Q then $q(t) = p_j(t)$. Also, note that if Q is accepted at time t then the seller s receives 'profit' $q(t) - \lambda_{s,j}$ and the buyer b receives profit $\lambda_{b,j} - q(t)$, where $q(t)$ is either $p_s(t)$ or $p_b(t)$.

A ZIP seller i raises its profit margin whenever Q was accepted and $p_i(t) \leq q(t)$. It lowers its margin only if it is still active and Q was an offer with $p_i(t) \geq q(t)$, or if Q was a bid that was accepted and $p_i(t) \geq q(t)$. Similarly, a ZIP buyer i raises its profit margin whenever Q was accepted and $p_i(t) \geq q(t)$, and it lowers its margin when it is active and either Q was a rejected bid with $p_i(t) \leq q(t)$ or Q was an accepted offer with $p_i(t) \leq q(t)$.

The quantitative issue of by how much the profit coefficient $\mu_i(t)$ should be altered is addressed by using a simple machine-learning rule, *Widrow-Hoff with momentum*, which also underlies back-propagation learning in neural networks [39]. Briefly, this adjusts the actual output of a system toward some *target* output value, at a speed determined by a learning rate β, and with a simple 'memory' or 'momentum' parameter γ. In each ZIP trader the target value $\tau_i(t)$ is a function of $q(t)$: $\tau_i(t) = \mathcal{R}_i(t)q(t) + \mathcal{A}_i(t)$ where $\mathcal{R}_i(t)$ and $\mathcal{A}_i(t)$ are stochastic functions that return uniformly distributed real values generated from IID distributions for each trader i. The intention is that $\mathcal{R}_i$ is a small *relative* perturbation of the last quote price ($|\mathcal{R}_i(t)| \simeq 1.0; \forall i \forall t$), and $\mathcal{A}_i$ is a small *absolute* perturbation ($|\mathcal{A}_i(t)| \simeq 0.0; \forall i \forall t$). Using $\mathcal{U}(c_l, c_u)$ to denote a uniform distribution over the real-valued range $[c_l, c_u]$, for buyers the $\mathcal{R}_i$ values are given by $\mathcal{U}(1.0-c_r, 1.0)$ and for sellers the $\mathcal{R}_i$ values are given by $\mathcal{U}(1.0, 1.0+c_r)$. Similarly, the $\mathcal{A}_i$ values for buyers are given by $\mathcal{U}(-c_a, 0.0)$ and for sellers by $\mathcal{U}(0.0, c_a)$. Thus, the distributions of the $\mathcal{R}_i$ and $\mathcal{A}_i$ values for the entire ZIP-trader market are determined by the two parameters c_r and c_a.

Each trader i uses its current value of $\tau_i(t)$ in combination with β_i and γ_i to adjust its profit coefficient $\mu_i(t)$. The profit-coefficient update rule for trader i is:

$$\mu_i(t+1) = (p_i(t) + \Gamma_i(t))/\lambda_{i,j} - 1$$

where: $\Gamma_i(t > 0) = \gamma_i \Gamma_i(t-1) + (1-\gamma_i)\beta_i(\tau_i(t) - p_i(t))$ and $\Gamma_i(0) = 0 : \forall i$. For further details of how learning is implemented in ZIP traders, see Cliff & Bruten [10, 14].

Thus, adaptation in each trader i is governed by three real-valued parameters: learning rate β_i, momentum γ_i, and initial profit coefficient $\mu_i(0)$. Each trader's values for these parameters are assigned at initialization, using uniform distributions: for all traders, β_i is assigned a value generated at random from $\mathcal{U}(\beta_b, \beta_b + \beta_\Delta)$; and γ_i is assigned a value from $\mathcal{U}(\gamma_b, \gamma_b + \gamma_\Delta)$.[3] For sellers, the initial profit coefficients $\mu_i(0)$ are assigned from $\mathcal{U}(\mu_b, \mu_h)$, and for buyers the $\mu_i(0)$ values are assigned from $\mathcal{U}(-\mu_h, -\mu_b)$, where $\mu_h = \mu_b + \mu_\Delta$.

Therefore, the adaptation of any ZIP-trader market is determined by eight real-valued parameters: the 3 pairs of bounds on the distribution of parameters for the individual agents (i.e., $\beta_b, \beta_\Delta, \gamma_b, \gamma_\Delta, \mu_b$, and μ_Δ), and the two parameters c_r and c_a that define the distributions of stochastic perturbations used in calculating each trader's target price. Clearly, any particular choice of values for these eight parameters can be represented as a vector V:

$$V = [\beta_b, \beta_\Delta, \gamma_b, \gamma_\Delta, \mu_b, \mu_\Delta, c_r, c_a] \in \Re^8$$

which corresponds to a single point in the 8-dimensional space of possible parameter values.

In previous publications, Cliff & Bruten demonstrated the effectiveness of the ZIP adaptation strategy in a variety of CDA markets [10, 14, 12] and also in simple one-sided auction models of retail markets [15]. In these publications, the dynamics of markets populated by ZIP traders were explored using various simple experimental supply and demand curves similar to those used by Smith [42] or Gode & Sunder [23]. Cliff & Bruten [10, 12] presented results illustrating ZIP traders operating successfully in CDA markets where Gode & Sunder's [23] zero-intelligence traders fail. These examples include markets where there are asymmetric supply and demand functions and imbalances between the number of buyers and sellers. Cliff & Bruten [10, 16] also showed that ZIP-trader markets can respond well to shock changes in supply and demand and that they can operate successfully in non-CDA markets, such as experimental 'retail' markets. Thus, in experimental markets such as those used by Smith or Gode & Sunder, the results from ZIP traders are very similar to those from Smith's [42] human subjects: a point explored in detail in [10]. Although Cliff & Bruten varied the details of the market environments, in all of their experiments the same vector of parameter values was used, denoted here by V_{cb}:

$$V_{cb} = [0.10, 0.40, 0.00, 0.10, 0.05, 0.30, 0.05, 0.05]$$

The values of the elements of V_{cb} were chosen using "educated guesses" followed by some trial-and-error experimentation to fine-tune the performance.

To demonstrate the typical performance of a CDA market of ZIP traders operating with parameters set by V_{cb}, Fig. 1 shows market supply and demand for 22 ZIP traders: there are 11 buyers and

[3] The b subscript in β_b, γ_b, and μ_b denotes the *baseline* value.

11 sellers, each with the right to engage in one transaction. The limit prices for both the buyers and the sellers range from $0.75 to $3.25 in steps of $0.25: the supply and demand curves induced by this distribution of limit prices intersects at an equilibrium price of $2.00. Fig. 2 shows the transaction-price time series for one CDA experiment in the market of Fig. 1. The experiment lasts for six trading periods or 'days'. Initial transaction prices occur at off-equilibrium prices but there is clear convergence to equilibrium over the course of this experiment. This convergence is quantified in Fig. 3, where summary statistics for ZIP performance in the market of Fig. 1 are shown: in each experiment, Smith's [42] α measure of convergence[4] is calculated for each day; Fig. 3 shows the mean and standard deviation of α in each day over 50 such experiments.

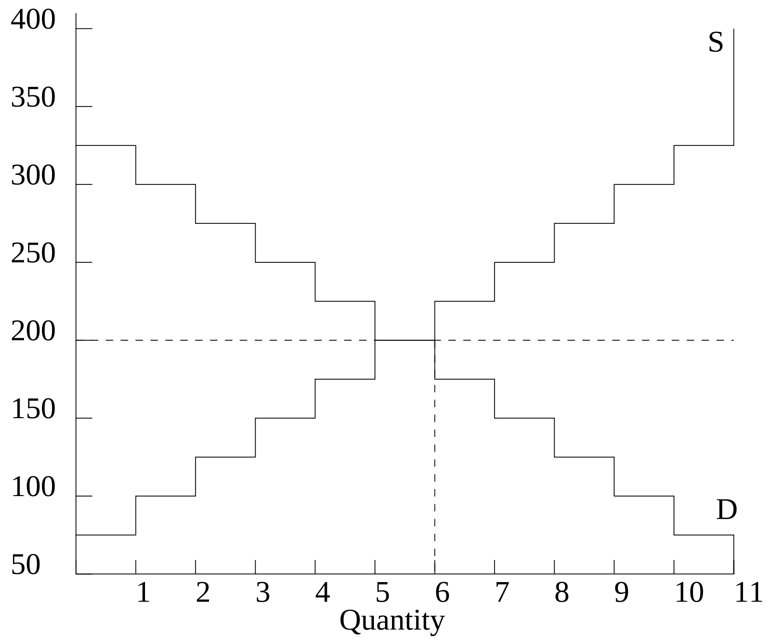

Figure 1: Market supply curve S and demand curve D, for 11 buyers and 11 sellers. Vertical axis is price ($0.50 to $4.00); equilibrium price is $2.00.

In developing and testing the ZIP traders, values for the elements of V_{cb} that gave good market performance were chosen by hand. This manual process is potentially labor-intensive: it is possible that many combinations of parameter values need to be evaluated for any given market situation before a satisfactory set of values is discovered. Furthermore, whenever the market alters, this hand-optimization process may need to be repeated. Feasible alterations to the market include significant shifts in the supply and demand curves, or changes in the number of buyers or sellers in the market. Thus, it would be highly desirable if the ZIP parameter-optimization process could be automated. There are many possible optimization techniques that could be employed. In the next section, results from the first attempt at optimizing via a genetic algorithm (GA) are described. GAs have been used for parameter optimization in a wide variety of problem domains [25]. The usage described below is reminiscent of some work in evolving parameters affecting back-propagation learning in artificial neural networks (e.g., [37, 2]), but the only connection between ZIP

[4]Smith's α measure of convergence for a trading period is given by σ_0/P_0 where P_0 is the theoretical equilibrium price and σ_0 is the root mean square of the difference between the transaction prices and P_0 during that trading period.

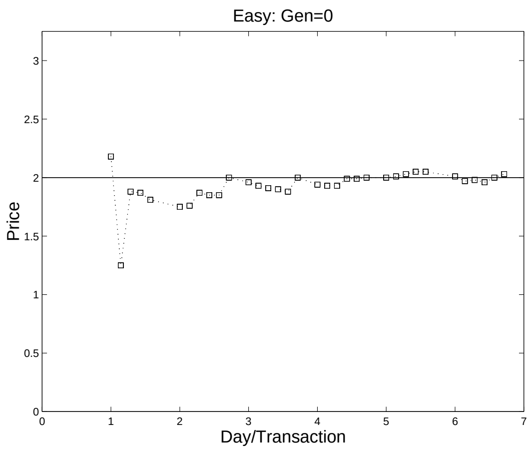

Figure 2: Transaction-price time series from one 6-day CDA market experiment with parameters set by V_{cb}. Vertical axis is price (\$0.00 to \$3.25); horizontal axis is day. Horizontal line indicates equilibrium price of \$2.00.

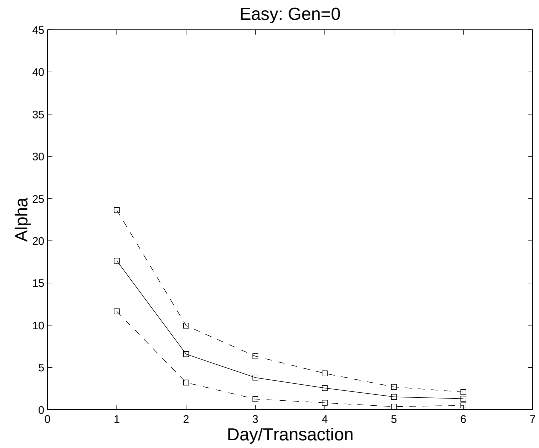

Figure 3: Mean (solid line) plus and minus one standard deviation (dashed lines) of α in each day over 50 6-day CDA market experiments with parameters set by V_{cb}. Vertical axis is Smith's α (0 to 45); horizontal axis is day.

traders and back-propagation networks is that both are based on use of the Widrow-Hoff learning rule with momentum. Future work is planned to explore alternative optimization strategies and compare their results with the GA approach described here.

3 Evolutionary Optimization

A standard GA (e.g. [25]) was used to optimize the parameters. An initial population of 30 'individuals' was randomly created, where each individual was a 'genotype' consisting of the 8 real values of a V vector. Each individual genotype V_i was evaluated by monitoring the market dynamics of groups of ZIP traders with parameters set by V_i, and the result of this monitoring was used to assign V_i a 'fitness' score. Once the fitness of all individuals had been evaluated, a new population of 30 was created via 'reproduction' where a selection process ensured that fitter individuals were more likely to reproduce. The old population would then be discarded and the new population would be evaluated to assign fitnesses that determine likelihood of reproduction. Each cycle of evaluating the current population and 'breeding' a new population from the fitter individuals is referred to as one 'generation'. In all the experiments discussed below, the GA was ended after 200 generations.

For each new individual that was created in the breeding phase, a simple rank-based tournament selection process was used: three distinct individuals were randomly selected from the old (evaluated) population, and the fittest two of these were identified as the 'parents' of the new individual. Let V_{mom} denote the fitter of the two parents and let V_{dad} denote the other parent. Also, let V_{kid} denote the new individual. The 'breeding' process started by copying the value of the first element of V_{mom} into the first element of V_{kid}. A uniform random value $x \in [0.0, 1.0]$ was then generated. If x was less than some threshold T_x then the copying process 'crossed-over' so that the next element of V_{kid} would be copied from V_{dad}; otherwise, copying continued from V_{mom}. This process of copying from the current parent and then crossing-over to the other parent with probability T_x was repeated at each of the eight elements. Thus, it is possible that V_{kid} is an *asexual* copy of V_{mom}, or that the values of the elements of V_{kid} are a *sexual* mix of the 'genetic material' of the two parents. Clearly, the number of cross-overs from one parent to the other in the copying process follows a Poisson distribution dependent on T_x. This so-called *stochastic multi-point crossover* is a standard practice in GA applications. In all the experiments reported here, $T_x = 0.125$. When each element was copied from a parent to its child, the value on the element was also 'mutated' by adding a random real value generated from $\mathcal{U}(-0.05, +0.05)$, and clipping at 0.0 and 1.0 to ensure that $V_i \in [0, 1]^8$.

The standard practice of *elitism* was also employed: on each generation, one unadulterated copy of the best individual in the old population was copied into the new population, thereby helping ensure that the best individual found so far is always retained. However, because the overall ZIP-trader market is stochastic, the fitness evaluation process is nondeterministic, and hence if the same population was to be evaluated twice, there is no guarantee that the same individual would be identified as the elite both times.

The 'fitness' for each individual V_i genotype was calculated by monitoring price convergence in a series of 50 CDA market experiments, all with the supply and demand curves shown in Fig. 1. At the start of each experiment the parameter values represented on V_i were used to generate the β_i, γ_i, and $\mu_i(0)$ values for the 22 ZIP traders. The values of c_r and c_a used in calculating the traders' target prices in each experiment were also taken from V_i. Each market experiment lasted for six trading 'days'. At the end of each day d, Smith's α measure would be calculated and is denoted here by $\alpha(d)$. The score for V_i on experiment number e, denoted by $S(V_i, e)$, was given by a weighted sum of the six $\alpha(d)$ values: let w_d denote the weight on day d. In the experiments reported here, $w_1 = 1.75, w_2 = 1.5, w_3 = 1.25$, and w_4, w_5, and w_6 were all equal to 1.0. These weights place a greater emphasis on the early trading days, when the ZIP traders are undergoing their initial adaptation to the market. The 'fitness' evaluation function for V_i, denoted $F(V_i)$ was then calculated as the arithmetic mean of $S(V_i, e)$ over n=50 experiments:

$$F(V_i) = \frac{1}{n}\sum_{e=1}^{n} S(V_i, e) = \frac{1}{n}\sum_{e=1}^{n}\frac{1}{6}\sum_{d=1}^{6} w_d \alpha(d)$$

Note that, because lower values of α correspond to better transaction-price convergence, lower values of F are more 'fit', and an optimal score is $F(V_i) = 0.0$. Thus the intention of the GA is to *minimize* this measure (i.e., the closer an individual's score is to 0.0, the more 'fit' it is).

Before results are shown, one more detail needs to be discussed: the choice of bounds on the distributions that generate the initial random population of V_i individuals. These determine the initial conditions of the evolutionary search, and thus could have a significant effect on the success of the search.[5] In the following section, results will be shown from experiments where the initial conditions are classified into three types:

- **Easy:** $V_i = V_{cb}; \forall i$.
- **Zero:** $V_i = [0, 0, 0, 0, 0, 0, 0, 0]; \forall i$.
- **Hard:** $V_i \in [0.75, \mathcal{U}_\Delta, 0.75, \mathcal{U}_\Delta, 0.75, \mathcal{U}_\Delta, \mathcal{U}_c, \mathcal{U}_c]; \forall i$
  where $\mathcal{U}_\Delta = \mathcal{U}(0.00, 0.25)$ and $\mathcal{U}_c = \mathcal{U}(0.75, 1.00)$.

The 'easy' type can be viewed as a test of whether the evolutionary optimization process can find *improvements* on the fitness of V_{cb}: effectively, the 'easy' experiments explore the possibility

[5]The values of any element of V in the initial population can be constrained to some constant c by generating them from $\mathcal{U}(c, c)$.

that, with further trial-and-error settings of the parameters, even better performance could be obtained from the ZIP traders.

The 'zero' experiments can be viewed as a test of whether the evolutionary process can find good sets of parameters from initial conditions that are intuitively suboptimal: all parameters set to zero corresponds to having *no adaptation at all* in the markets; yet having *some* adaptation seems intuitively to be a better approach – the GA offers a less subjective test of this belief. Thus, the 'zero' experiments can be characterized as a test of the hypothesis that some learning or adaptation is better than none at all.

However, it is clear that V_{cb} is quite close to an all-zero vector (in terms of Euclidian distance at least). So, starting with a population of all-zero vectors may not present much of a challenge to the search process. For this reason, the 'hard' experiments commence with a set of parameter values that are all very much higher than would be chosen on the basis of intuition or experience. Having high learning rates, high momentum coefficients, high initial profit margins, and large perturbations when computing the target price, are all likely to give rise to poor market dynamics. But, it is not clear *a priori* that the evolutionary process will be able to guide the population from this area of parameter space with poor dynamics to one that gives better dynamics: the possibility that the system will be trapped in local minima cannot be ruled out.

Illustrative end-results from these three types of evolutionary optimization experiment are presented in the next section. Following that, the evolutionary trajectories of the population through parameter space are discussed.

4 End-Results

Fig. 4 shows results from one 'easy' experiment: it is a log-log plot of the fitnesses of the 30 individuals in the population at each of the 200 generations. As can be seen, there is an initial small but significant increase in performance: the elite score falls from 7.276 in the initial population down to about 80% of that initial elite score by generation 50, from which point onwards no significant improvement in scores is discernible. In the last 20 generations of this experiment the lowest elite score is 5.289, the highest is 6.549, the mean is 6.092 and the standard deviation is 0.312. Fig. 5 shows a time-series of transaction-prices from the elite individual in generation 200, and Fig. 6 shows a summary of the α values from 50 such experiments. These results should be compared to those shown in Figs 2 and 3, where results from the V_{cb} parameter set (as used in the initial generation of the 'easy' experiments) are illustrated. While the $\alpha(1)$ values are similar, the generation-200 individual has better values for $\alpha(d > 1)$. Thus, the evolutionary search has discovered parameter settings that give better convergence dynamics than those given by V_{cb} in the market of Fig. 1. The final elite individual in this experiment has the vector V_{easy}:

$$V_{\text{easy}} = [0.438, 0.456, 0.095, 0.171, 0.209, 0.187, 0.000, 0.071]$$

To give an indication of how representative this is of the elites in the final 20 generations of the experiment, let $V_{\text{easy}:\mu}$ denote the vector of arithmetic means of the elements of the last 20 elite genotypes, and let $V_{\text{easy}:\sigma}$ denote the corresponding standard deviations. These vectors are:

$$V_{\text{easy}:\mu} = [0.487, 0.343, 0.073, 0.141, 0.161, 0.158, 0.006, 0.050]$$

$$V_{\text{easy}:\sigma} = [0.056, 0.056, 0.040, 0.064, 0.036, 0.046, 0.009, 0.023]$$

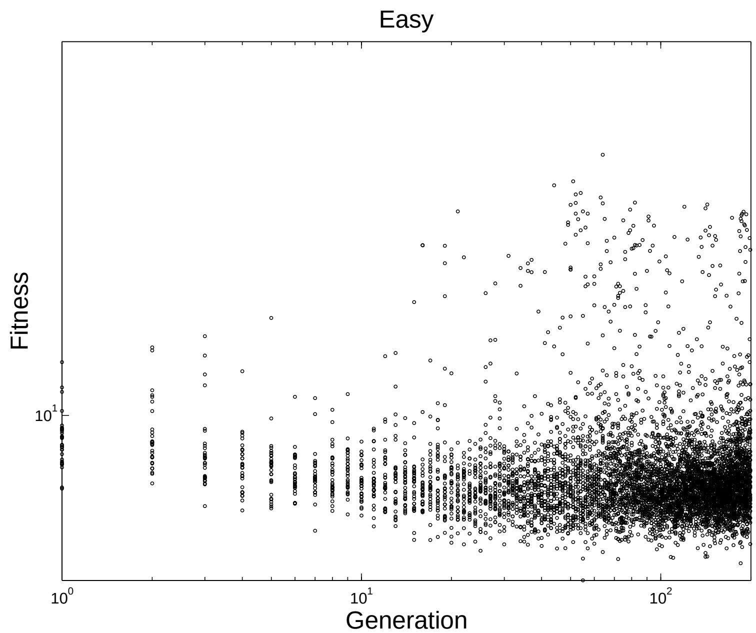

Figure 4: Log-log plot of the fitness (vertical axis; range is [0,50]) of each of the 30 individuals in the population for one experiment over 200 generations (horizontal axis) from 'easy' initial conditions.

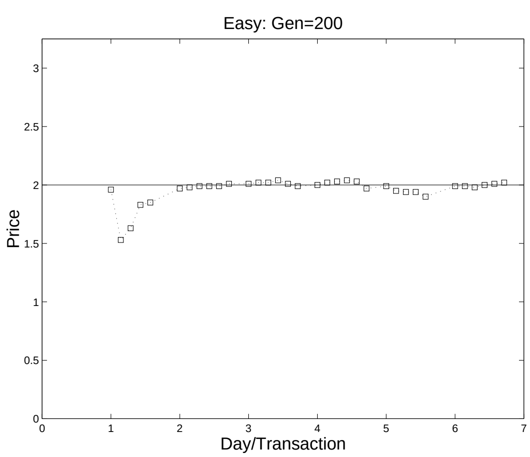

Figure 5: Transaction-price time series from one experiment with the elite individual in the final generation from 'easy' initial conditions. Axis ranges same as in Fig. 2.

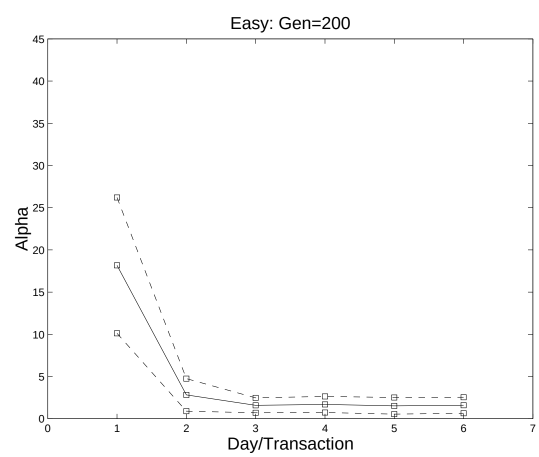

Figure 6: Mean and standard deviation of α in each day over 50 experiments for the elite individual in the final generation from 'easy' initial conditions. Axis ranges same as in Fig. 3.

Fig. 7 shows the population fitnesses in one 'zero' experiment, using the same format as Fig. 4. Again, most of the improvement in F scores comes over the first 50 generations, but the improvement is much more dramatic in this case: the best F value in the initial population is 44.236, confirming the intuitive assumption that some adaptation is (much) better than no adaptation at all. The performance of members of the initial population is illustrated in the transaction-price time series of Fig. 8 and the $\alpha(d)$ statistics of Fig. 9. The corresponding time-series and $\alpha(d)$ data for the elite individual after 200 generations are shown in Figs 10 and 11. In the last 20 generations of this experiment, the lowest elite score was 5.345, the highest was 6.223, the mean was 5.916, and the standard deviation was 0.218. The final elite individual has the vector V_{zero}:

$$V_{\mathrm{zero}} = [0.621, 0.000, 0.001, 0.277, 0.214, 0.063, 0.031, 0.010]$$

and the mean $V_{\mathrm{zero}:\mu}$ and standard deviation $V_{\mathrm{zero}:\sigma}$ vectors for the elites of the final 20 generations are:

$$V_{\mathrm{zero}:\mu} = [0.539, 0.026, 0.049, 0.339, 0.204, 0.053, 0.014, 0.047]$$

$$V_{\mathrm{zero}:\sigma} = [0.066, 0.021, 0.049, 0.090, 0.030, 0.030, 0.012, 0.032]$$

Fig. 12 shows the population fitnesses in one 'hard' experiment, using the same format as for Figs 4 and 7. In comparison to those figures, the improvement in F scores takes much longer: the scores do not level off until after approximately 100 generations. The performance of the elite of the initial population is illustrated in the transaction-price time series of Fig. 13 and the $\alpha(d)$ statistics of Fig. 14: this individual had a score of 21.591. The time-series and $\alpha(d)$ data for the elite individual after 200 generations are shown in Figs 15 and 16. In the last 20 generations of this experiment, the lowest elite score was 5.760, the highest was 6.637, the mean was 6.209, and the standard deviation was 0.229. The final elite individual vector was V_{hard}:

$$V_{\mathrm{hard}} = [0.361, 0.206, 0.000, 0.444, 0.191, 0.075, 0.000, 0.118]$$

and the mean $V_{\mathrm{hard}:\mu}$ and standard deviation $V_{\mathrm{hard}:\sigma}$ vectors for the elites of the last 20 generations are:

$$V_{\mathrm{hard}:\mu} = [0.462, 0.173, 0.040, 0.481, 0.178, 0.032, 0.008, 0.070]$$

$$V_{\mathrm{hard}:\sigma} = [0.070, 0.048, 0.037, 0.043, 0.025, 0.026, 0.010, 0.033]$$

Comparing Figs 6, 11, and 16, it is clear that the final elite genotypes in each of the three experiments give very similar market-convergence performance. In all three experiments, the elite vectors (both the final elite individual, and the mean elite genotype of generations 180–200) show regular departures from V_{cb}. In comparison to V_{cb}, all three evolved elites have values of β_b much higher than the value in V_{cb}, and the value of c_r in all the evolved elites is very close to zero. Yet the three experiments have not evolved to identical (or statistically indistinguishable) values. Examining the three mean elite vectors, it is clear that while the values for $\beta_b, \gamma_b, \mu_b, c_r$, and c_a are all roughly the same (within approximately one standard deviation of each other), the values for $\beta_\Delta, \gamma_\Delta$, and μ_Δ vary significantly. Yet these variations have little impact on the observable market-convergence performance of the elites.

Naturally, the most important point to note from these results is that in all three initial-condition cases, the simple GA can find parameter vectors that give good α convergence in the market of Fig. 1.

Nevertheless, there are two other issues that deserve consideration. The first is that in all three experiments the evolved value of c_r was either zero or very close to zero. The implication

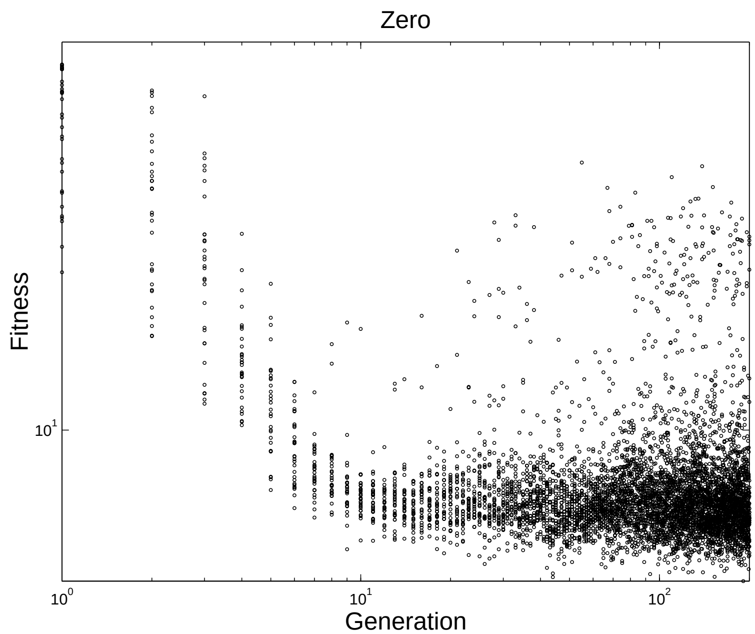

Figure 7: Log-log plot of the fitness of each of the 30 individuals in the population for one experiment over 200 generations from 'zero' initial conditions. Axis ranges same as in Fig. 4.

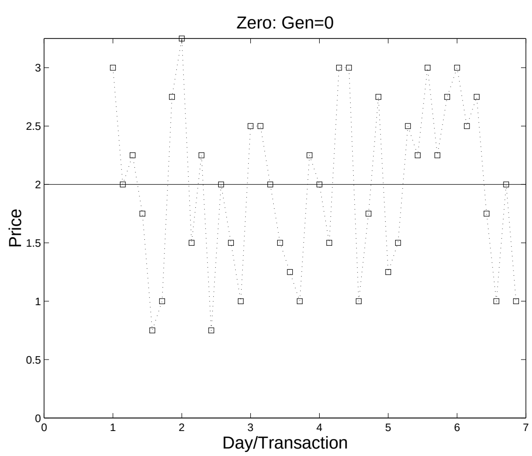

Figure 8: Transaction-price time series from one experiment with the elite individual in the first generation from 'zero' initial conditions. Axis ranges same as in Fig. 2.

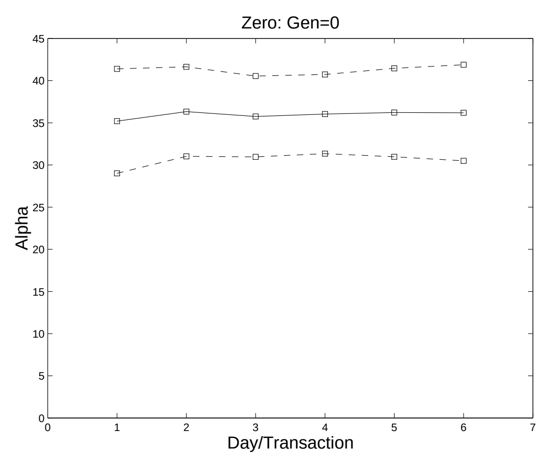

Figure 9: Mean and standard deviation of α in each day over 50 experiments for the elite individual in the first generation from 'zero' initial conditions. Axis ranges same as in Fig. 3.

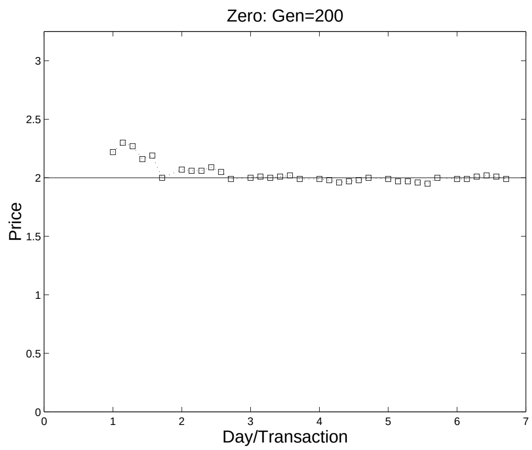

Figure 10: Transaction-price time series from one experiment with the elite individual in the final generation from 'zero' initial conditions. Axis ranges same as in Fig. 2.

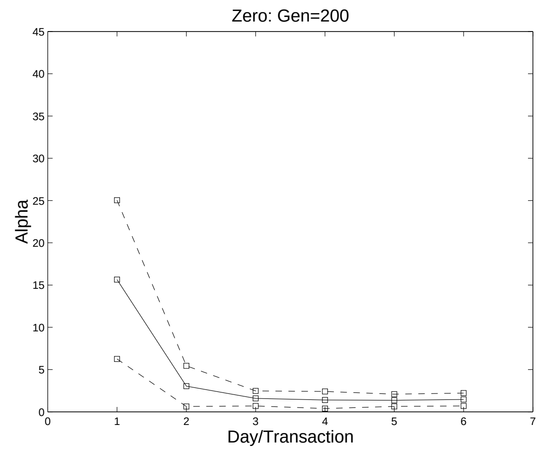

Figure 11: Mean and standard deviation of α in each day over 50 experiments for the elite individual in the final generation from 'zero' initial conditions. Axis ranges same as in Fig. 3.

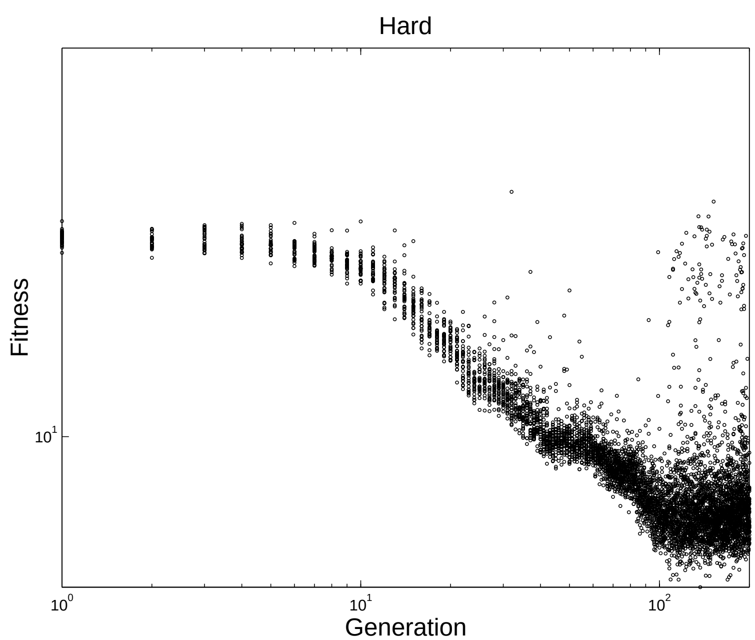

Figure 12: Log-log plot of the fitness of each of the 30 individuals in the population over 200 generations from 'hard' initial conditions. Axis ranges same as in Fig. 4.

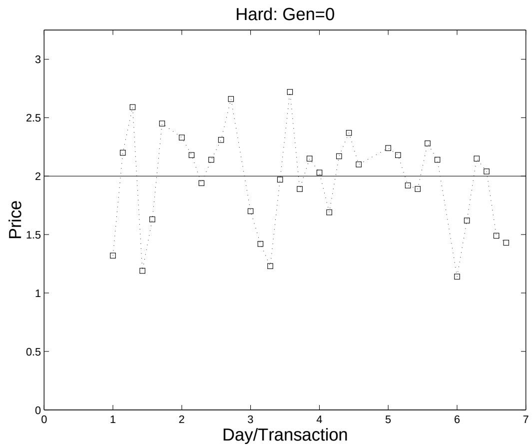

Figure 13: Transaction-price time series from one experiment with the elite individual in the first generation from 'hard' initial conditions. Axis ranges same as in Fig. 2.

Figure 14: Mean and standard deviation of α in each day over 50 experiments for the elite individual in the first generation from 'hard' initial conditions. Axis ranges same as in Fig. 3.

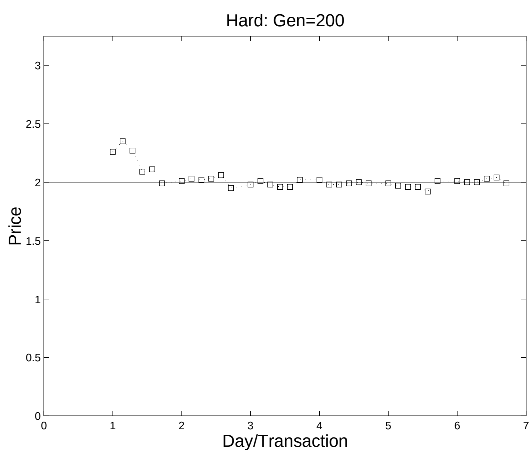

Figure 15: Transaction-price time series from one experiment with the elite individual in the final generation from 'hard' initial conditions. Axis ranges same as in Fig. 2.

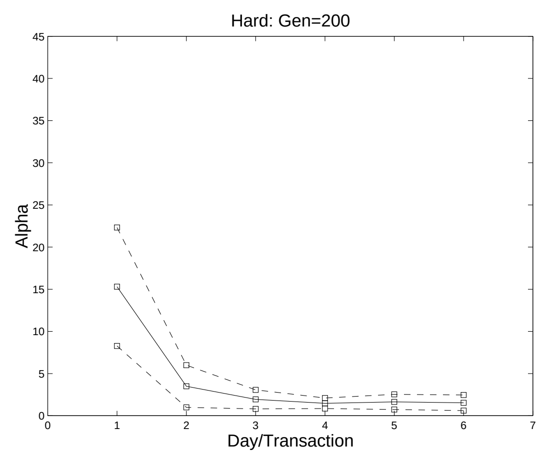

Figure 16: Mean and standard deviation of α in each day over 50 experiments for the elite individual in the final generation from 'hard' initial conditions. Axis ranges same as in Fig. 3.

of this could be that the c_r parameter is 'harmful', in the sense that non-zero values of c_r yield less fitness than zero values, and setting $c_r = 0.0$ has the effect of disabling relative perturbations of the quote price $q(t)$ when calculating the trader's target price $\tau_i(t)$. Higher fitness appear to result from having no such perturbations. Although there was a design rationale for having these relative perturbations, the results of these three evolutionary experiments indicates that better performance is given when the relative perturbations are *prevented* from occurring. If this is true, it may only be so for the one particular market supply and demand used in these experiments (Fig. 1): in other markets, the 'harmful' parameters may usefully take nonzero values, while other parameters may be 'harmful'. The second issue is that the variation in the final evolved values of some of the parameters may indicate that they are 'redundant', in the sense that comparatively large variations in the values of these parameters has little impact on the performance of the system so long as the non-redundant parameters are set to appropriate values. If this is the case, then the variation in the evolved values of these parameters could simply be the result of the evolutionary process allowing a 'random walk' through the range of values for each 'redundant' parameter. The apparent variation is then simply the result of arbitrarily halting the random walk after 200 generations. More formally, the 'redundant' variables create a set of numerically distinct vectors (i.e., a set of differing genotypes) that constitute a 'plateau' in the fitness landscape, in the sense that differences between the members of the set are selectively neutral.

While such arguments sound plausible, it is difficult to judge their worth because, thus far, only the end-results of the evolutionary process have been examined; and this has only been done once for each type of initial-conditions. To gain a better understanding of how the evolutionary process proceeds, it is necessary to explore the dynamics of the GA in more detail, performing multiple repetitions of experiments from the same set of initial conditions. The next section presents such an exploration. Although experiments with any of the three initial-condition types could have been explored, the discussion here centers on the dynamics of the 'hard' experiments: as is clear from the results presented above, the 'hard' experiments appear to take longer to converge on a solution, and these longer transients offer the opportunity for more detailed monitoring of evolutionary dynamics.

5 Evolutionary Trajectories

Just as Figs 4, 7, and 12 show the fitness of each individual in the population over 200 generations of evolution, so it is possible to show the values of each of the 8 elements of the genotypes for each member of the population at each generation. Figs 17 to 24 show the 200-generation evolutionary trajectory of the population through the 8-dimensional parameter space during the 'hard' experiment of Fig. 12.

In the discussion of Fig. 12 it was noted that the improvement in elite scores takes place over roughly the first 100 generations, after which no further improvement is discernible and the elite scores stabilize. Qualitatively, this is reflected in all the parameter-value trajectories: data from the first half of the experiment in each figure shows, in most cases, significant shifts in the values maintained by the population; in the second half, most of the values are roughly constant, or there is drifting variation that has no effect on the elite fitness scores.

The data for $\gamma_b, \mu_b, \mu_\Delta, c_r$, and c_a (Figs. 19 and 21 to 24 respectively) show clear stasis in the second half of the experiment, while the data for β_b, β_Δ, and γ_Δ (Figs 17, 18, and 20 respectively) are more reminiscent of random walks in the second half of the experiment. Certainly, the observable changes in the population's distribution of values for these parameters has no discernible effect on the fitness scores of the population: thus these three variables may be considered 'redundant' in the sense discussed earlier.

The data for the first 100 generations in Figs 17 to 24 shows some intriguing differences. While the figures for μ_b, c_r, and c_a all clearly exhibit a roughly monotonic approach from the initial conditions to the range of values that the system settles on in the second half of the experiment, and the population's values for μ_Δ and γ_Δ stay in roughly the same range as the initial conditions for the first 100 generations, the data for β_b, γ_b, and μ_Δ all show pronounced nonmonotic trajectories. The values for β_b rapidly fall from the initial conditions around 0.75 to levels around 0.1 after 20 generations; they then increase more gradually to values around 0.5 by generation 100. The values for γ_b show a similar decline to near-zero, albeit less fast (taking about 60 generations), which is followed by about 20 generations of increase toward 0.5, a further 20 generations where the distribution of values 'splits', indicating a bimodal population, before returning back toward the

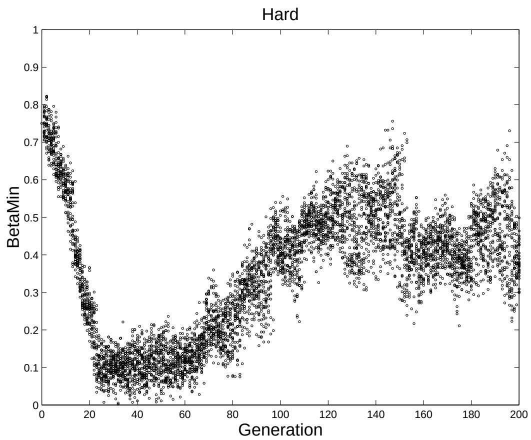

Figure 17: Values of β_b (vertical axis: linear [0.0,1.0]), for each of the 30 individuals in the population over 200 generations (horizontal axis: linear [0,200]).

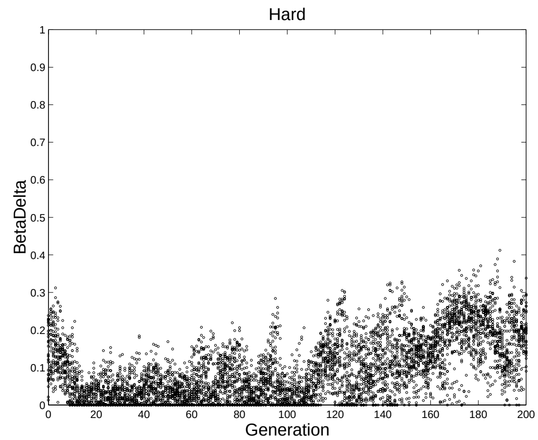

Figure 18: Values of β_Δ (vertical axis: linear [0.0,1.0]), for each of the 30 individuals in the population over 200 generations (horizontal axis: linear [0,200]).

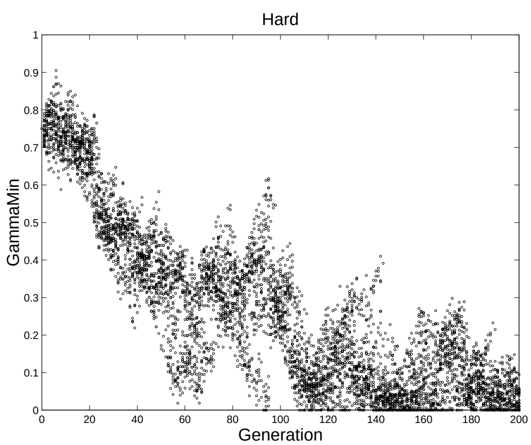

Figure 19: Values of γ_b (vertical axis: linear [0.0,1.0]), for each of the 30 individuals in the population over 200 generations (horizontal axis: linear [0,200]).

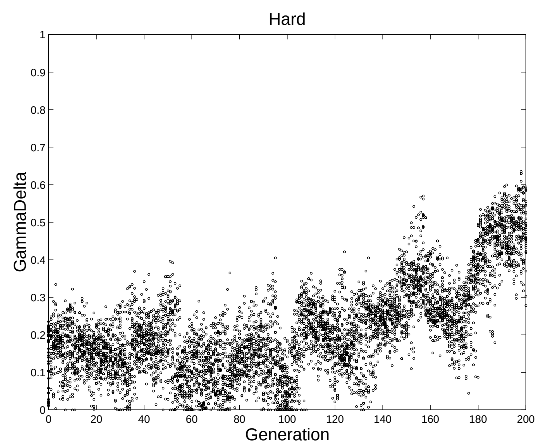

Figure 20: Values of γ_Δ (vertical axis: linear [0.0,1.0]), for each of the 30 individuals in the population over 200 generations (horizontal axis: linear [0,200]).

zero level. The values for μ_Δ start low, but increase to around 0.4–0.6 over the first 40 generations, falling down to $\simeq$0.1 over generations 40–60.

As was noted in [11], the trader learning rates β_i appear to be a crucial parameter. Initially, reducing the β_i values has the greatest beneficial effect on fitness: β_i, which essentially controls the 'gain' in the ZIP adaptation process, is set at such a highly suboptimal value that ZIP-trader adaptation causes more problems than it solves. Hence greater fitness is found in genotypes that 'turn down' the learning rate, thereby reducing these ill effects. More formally: initially, with all parameters having high values, there is a negative correlation between fitness and learning rates β_i.

However, once the distribution of β_i values is sufficiently low, the effects on fitness of variations in other parameters becomes more pronounced. This is seen most clearly in Fig. 21, where there is a

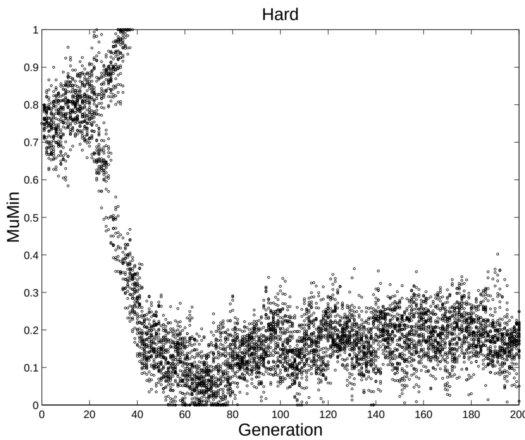

Figure 21: Values of μ_b (vertical axis: linear [0.0,1.0]), for each of the 30 individuals in the population over 200 generations (horizontal axis: linear [0,200]).

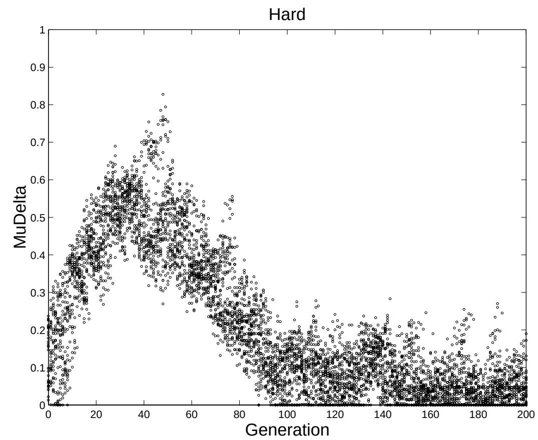

Figure 22: Values of μ_Δ (vertical axis: linear [0.0,1.0]), for each of the 30 individuals in the population over 200 generations (horizontal axis: linear [0,200]).

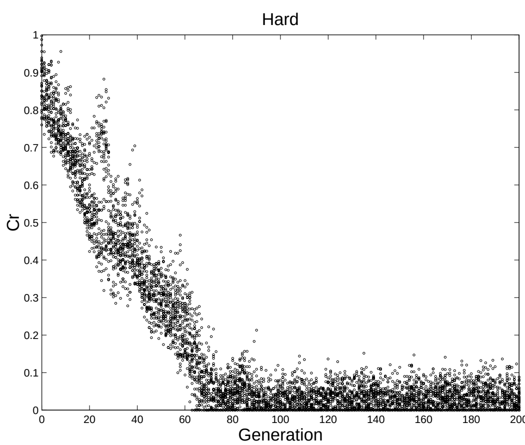

Figure 23: Values of c_r (vertical axis: linear [0.0,1.0]), for each of the 30 individuals in the population over 200 generations (horizontal axis: linear [0,200]).

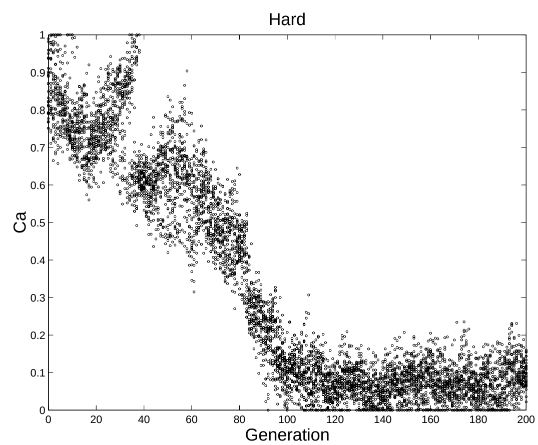

Figure 24: Values of c_a (vertical axis: linear [0.0,1.0]), for each of the 30 individuals in the population over 200 generations (horizontal axis: linear [0,200]).

sharp downwards shift, almost a jump discontinuity, in the population distribution of values for μ_b. This occurs around generations 25–30, i.e. shortly after the population values of β_b have stabilized at low values. Similar interactions between variables occur subsequently in the evolutionary process. As the values of all the variables evolve into ranges that yield effective market-adaptation in the traders, so traders with higher learning rates will adapt even faster, and so thus the initial negative correlation between β_i and fitness is *reversed*: now genotypes coding for higher distributions of β_i are likely to yield higher fitness, so there is a positive correlation. The data for β_b over generations 60–120 clearly shows the population evolving into regions of parameter space with higher β_b values.

Although Figs 17 to 24 faithfully illustrate the evolutionary trajectory of the population through the 8-d parameter space, they obscure the evolutionary trajectory of the parameter values encoded on the elite genotype. While it would be informative to present another sequence of graphs, one for each of the eight parameters, showing the values on the elite genotype at each generation, it is possible to represent such data in a more compact form. Because the parameters $\beta_b, \beta_\Delta, \gamma_b, \gamma_\Delta, \mu_b$, and μ_Δ are bounds in a 3-d $(\beta_i, \gamma_i, \mu(0)_i)$ space, the values of these six parameters define a rectangular solid subspace or 'box' in the 3-d $(\beta_i, \gamma_i, \mu(0)_i)$ space: the values of β_b, γ_b, and μ_b define one corner of a box, while the values of $\beta_\Delta, \gamma_\Delta$, and μ_Δ define the width, depth, and height of the box. Thus, for any one genotype, the first six parameters on a V_i genotype can be visualized as a box in 3-space. Figure 25 shows one $(\beta_i, \gamma_i, \mu(0)_i)$ box for the elite genotype in the population at each generation in the 'hard' experiment illustrated in Fig. 12. The evolutionary trajectory of the elite parameter values is fairly clear from comparing Fig. 25 to the data in Figs 17 to 24 (and even clearer when watching animations of the development of Fig 25).

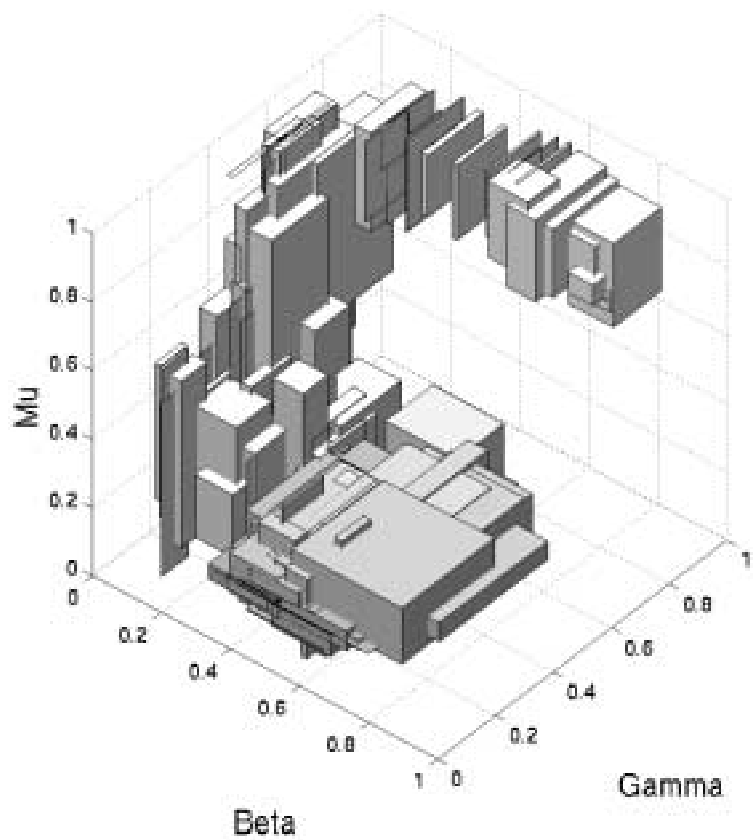

Figure 25: Evolutionary trajectory of elite genotype 'boxes' over 200 generations through $(\beta_i, \gamma_i, \mu(0)_i)$ space from 'hard' initial conditions. Boxes for individuals from earlier generations have lighter-shaded top-surfaces.

The discussion so far has established that the evolutionary trajectories of the population through parameter space can be visualized and that the vizualizations can inform on the interactions between the variables. The evolutionary trajectory can be monitored using quantitative techniques such as those recently developed by Mayley [35]. This can help establish when variation in elite parameter values is significant, and when it is the result of random drift. However, the discussion has concentrated on analyzing the trajectory given by one experiment. A natural question to ask is how representative this one experiment is. Determining the answer to that question requires data from multiple experiments where the initial conditions are identical and the only source of variation is in the (pseudo-)random processes governing generation of the initial population, and subsequent selection, crossover, and mutation. To explore this issue, 15 additional 'hard' experiments were run, differing only in the initial seed used by the random-number generator.

Fig. 26 shows the elite fitness scores from all 16 'hard' experiments, and clearly illustrates a bimodality in the results: although twelve of the sixteen experiments settled to give elite fitnesses near 6.0, the remaining four experiments settled with elite fitnesses around 8.0. The occurrence of 'bad' populations indicates the presence of local minima in the search space. More precisely, the average score of the last 20 elite individuals in the sixteen 'good' experiments was 6.066 and the standard deviation was 0.265 (N=12×20=240). The average score of the last 20 elite individuals in the four 'bad' experiments was 7.898, with a standard deviation of 0.260 (N=4×20=80). Let $V_{\text{hg}:\mu}$ represent the mean parameter vector from the N=240 'good' elite individuals, and $V_{\text{hg}:\sigma}$ represent the corresponding standard deviations; similarly let $V_{\text{hb}:\mu}$ and $V_{\text{hb}:\mu}$ represent the same statistics for the N=80 'bad' elite individuals. The elements of the vectors are:

$$V_{\text{hg}:\mu} = [0.543, 0.119, 0.177, 0.155, 0.177, 0.079, 0.010, 0.057]$$

$$V_{\text{hg}:\sigma} = [0.110, 0.100, 0.107, 0.140, 0.046, 0.066, 0.013, 0.030]$$

$$V_{\text{hb}:\mu} = [0.125, 0.021, 0.210, 0.247, 0.991, 0.982, 0.019, 0.905]$$

$$V_{\text{hb}:\sigma} = [0.023, 0.027, 0.122, 0.168, 0.015, 0.021, 0.028, 0.076]$$

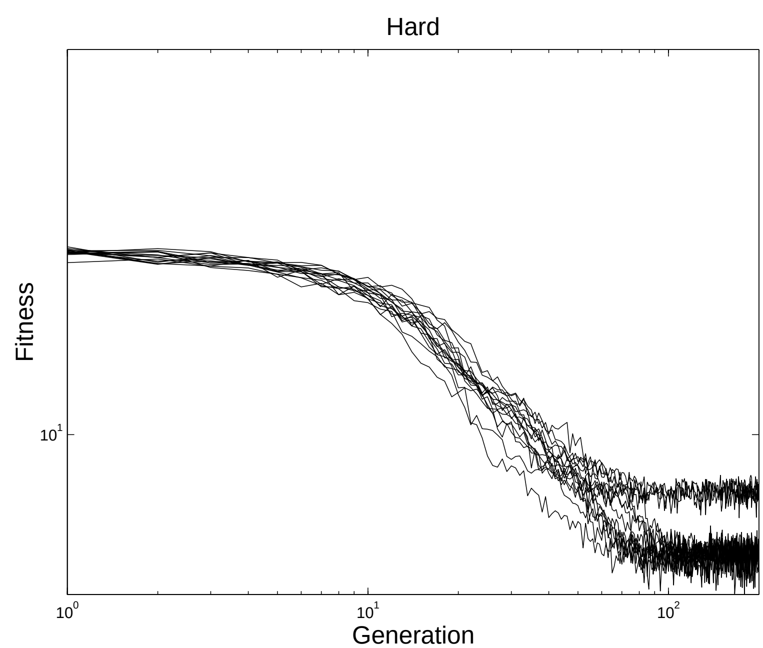

Figure 26: Scores of elite genotype over 200 generations in each of 16 'hard' experiments. Axis ranges same as in Fig. 4.

The primary differences between $V_{\text{hg}:\mu}$ and $V_{\text{hb}:\mu}$ are in the values of $\beta_b, \beta_\Delta, \mu_b, \mu_\Delta$, and c_a. Examining the evolutionary history of these differences shows that the 'bad' results occur when the population of parameter values follows a monotonic path through the parameter space. For

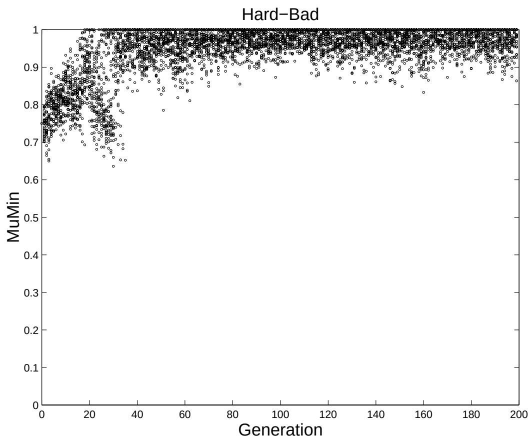

Figure 27: Values of μ_b for a 'bad' population.

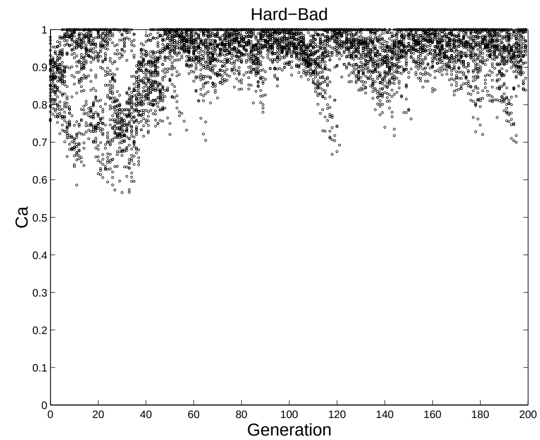

Figure 28: Values of c_a for a 'bad' population.

comparison, Figs 27 and 28 show the evolutionary trajectories of the population's values for μ_b and c_a in one of the 'hard' experiments that gave 'bad' results. In both cases, the first 20–40 generations do not appear significantly different from the corresponding 'good' data shown in Figs 21 and 24. However, the sudden drop in parameter values seen around generation 40 in the 'good' population does not occur in the 'bad' data. Fig. 29 shows the trajectory for the elite genotype 'boxes' in this 'bad' population. The early stages of this 'bad' trajectory are similar to those of the 'good' one shown in Fig. 25: both trajectories show a reduction in β values. However, the 'bad' population continues to remain in a region of very high μ and fairly high γ, rather than dropping down to the zone of low-μ low-γ where subsequent increases in β are possible. Colloquially, the problem with the 'bad' trajectory is not so much that it "takes a wrong turning" but rather that it fails to take the turning at all.

Finally, two notes of caution. First, note that care should be taken when discussing average genotypes such as $V_{\mathrm{hg}:\mu}$ and $V_{\mathrm{hb}:\mu}$: there is no guarantee that the average score of a set of genotypes accurately reflects the score of the average genotype. For example, if all elite genotypes are distributed evenly on the surface of a hypersphere, the average elite genotype will be at the origin of the hypersphere, where fitness may be very different.

Second, note that the 'easy', 'zero', and 'hard' initial-condition vector spaces were chosen for the reasons given in Section 3 (namely: to test whether the GA could improve on the hand-optimized V_{cb} values; to test whether any adaptation is better than none at all; and to test whether the GA can find good parameter-sets from manifestly poor start-conditions). A truly naive approach, ignoring the intended use of the parameter values, would be to randomly scatter the initial population throughout the parameter space. That is, to use initial vectors $V_i \in \mathcal{U}(0,1)^8 \in \Re^8; \forall i$. If the population is sufficiently large, then it is likely that at least one individual in the initial 'naive' population will be close to (or actually within) the zone of parameter space that the successful 'easy', 'zero' and 'hard' populations converge to. As the elite individual is always preserved into the next generation, the 'naive' experiment would reduce to watching material from the initial elite genotype spread through the population. This is of little explanatory interest, and so such results have not been presented here.

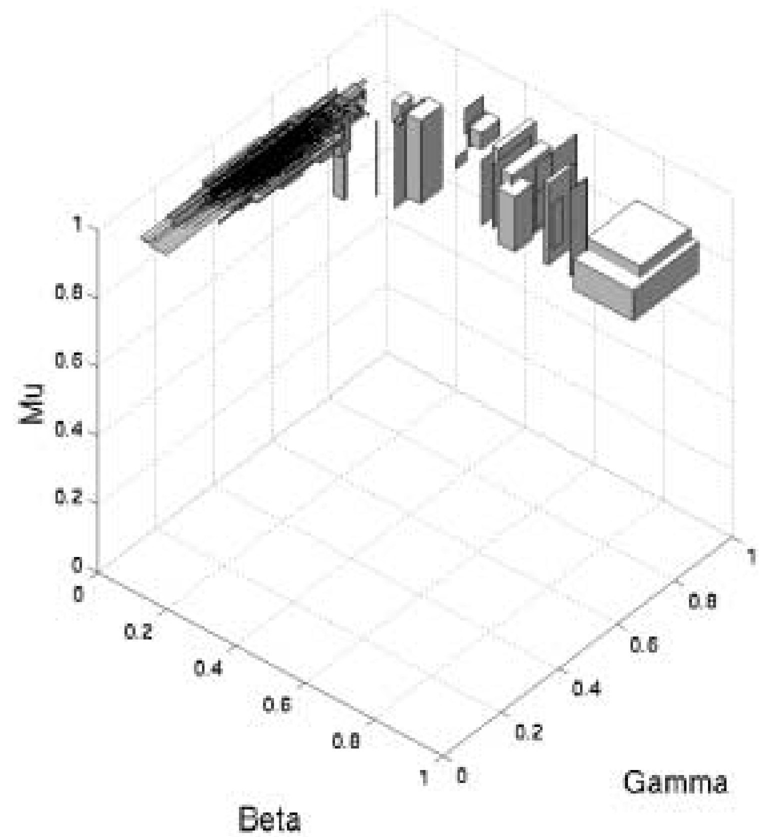

Figure 29: Evolutionary trajectory of elite genotype 'boxes' for a 'bad' population starting with 'hard' initial conditions.

6 Conclusion

This paper has demonstrated the use of a genetic algorithm in parameter-optimization for ZIP traders in a single CDA market with a fixed supply and demand schedule. The results indicate that the approach has merit.

Current research is directed at testing the capabilities of ZIP traders in more realistic and challenging environments, and extending them to the point where they can be used in market environments where time is continuous, trading is asynchronous, market supply and demand are dynamic, and information propagates with uncertainty and delays. For ZIP traders to operate in such environments, it is likely that additional parameters will have to be introduced. While manual optimization of low-dimensional parameter spaces such as the 8-dimensional space discussed here can be relatively straightforward, it is likely that manual techniques will grow unwieldy or impracticable as the number of parameters or dimensions increases. In such situations, automatic optimization is a highly desirable alternative, and the preliminary results presented here indicate that genetic algorithms are a promising approach.

The results presented here are for a simple fixed-demand, fixed-supply, discrete-time, single-commodity, unit-volume-per-transaction CDA market system, but by experimenting with simple systems such as this, likely difficulties can be revealed and possible solutions explored, with relatively low costs in terms of computer time required to evaluate alternative approaches, and increased ease of analysis. Results from this simple test-bed indicate that more complex artificial-trader systems (with tens, hundreds, or thousands of parameters) can be constructed with an increased reliance on (semi-)automatic optimization techniques.

References

[1] M. J. Anderson and S. Sunder. Professional traders and intuitive bayesians. *Organizational Behavior and Human Decision Processes*, 64(2):185–202, 1995.

[2] R. K. Belew, J. McInerney, and N. N. Schraudolph. Evolving networks: Using the genetic algorithm with connectionist learning. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, editors, *Artificial Life II: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, volume X of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 511–548, Redwood City, CA, 1992. Addison-Wesley.

[3] T. Bollerslev and I. Domowitz. Some effects of restricting the electronic order book in an automated trade execution system. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories, and Evidence*, pages 221–252. Addison-Wesley, New York, 1992.

[4] C. Camerer. Individual decision making. In J. H. Kagel and A. E. Roth, editors, *The Handbook of Experimental Economics*, pages 587–703. Princeton University Press, Princeton, NJ, 1995.

[5] T. Cason and D. Friedman. An empirical analysis of price formation in double auction markets. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories, and Evidence*, pages 253–284. Addison-Wesley, New York, 1992.

[6] T. Cason and D. Friedman. Price formation in double auction markets. *Journal of Economic Dynamics and Control*, 20(8):1307–1337, 1996.

[7] T. Cason and D. Friedman. Price formation in single call markets. *Econometrica*, 65(2):311–345, 1997.

[8] A. J. Clark. *Being There: Putting Brain, Body, and World Together Again.* MIT Press Bradford Books, Cambridge MA, 1997.

[9] S. H. Clearwater, editor. *Market-Based Control: A Paradigm for Distributed Resource Allocation.* World Scientific, Singapore, 1996.

[10] D. Cliff. Minimal-intelligence agents for bargaining behaviors in market-based environments. Technical Report HPL–97–91, Hewlett Packard Laboratories, Bristol, England, 1997.

[11] D. Cliff. Genetic optimization of adaptive trading agents for double-auction markets. In *Proceedings of the 1998 Conference on Computational Intelligence for Financial Engineering (CIFEr)*, pages 252–258. IEEE/IAFE/INFORMS, 1998. Preprint proceedings.

[12] D. Cliff and J. Bruten. Less than human: Simple adaptive trading agents for CDA markets. Technical Report HPL–97–155, Hewlett Packard Laboratories, Bristol, England, 1997. To be presented at CEFEES98, Cambridge UK, July 1998.

[13] D. Cliff and J. Bruten. More than zero intelligence needed for continuous double-auction markets. Technical Report HPL–97–157, Hewlett Packard Laboratories, Bristol, England, 1997. To be presented at CEFEES98, Cambridge UK, July 1998.

[14] D. Cliff and J. Bruten. Zero is not enough: On the lower limit of agent intelligence for continuous double auction markets. Technical Report HPL–97–141, Hewlett Packard Laboratories, Bristol, England, 1997. Submitted to *Journal of Political Economy*.

[15] D. Cliff and J. Bruten. Shop 'til you drop II: Collective adaptive behavior of simple autonomous trading agents in simulated retail markets. Technical Report HPL–98–59, Hewlett Packard Laboratories, Bristol, England, 1998.

[16] D. Cliff and J. Bruten. Simple bargaining agents for decentralized market-based control. Technical Report HPL–98–17, Hewlett Packard Laboratories, Bristol, England, 1998. To be presented at ESM98, Manchester UK, June 1998.

[17] J. Conlisk. Why bounded rationality? *Journal of Economic Literature*, 34(2):669–700, 1996.

[18] J. C. Cox and D. M. Grether. The preference-reversal phenomenon: Response-mode, markets, and incentives. *Economic Theory*, 7(3):381–405, 1996.

[19] D. D. Davis and C. A. Holt. *Experimental Economics*. Princeton University Press, Princeton, NJ, 1993.

[20] D. A. Evans. The role of markets in reducing expected utility violations. *Journal of Political Economy*, 105(3):622–636, 1997.

[21] D. Friedman. The double auction market institution: A survey. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories, and Evidence*, pages 3–26. Addison-Wesley, New York, 1992.

[22] D. Friedman and J. Rust, editors. *The Double Auction Market: Institutions, Theories, and Evidence*. Addison-Wesley, New York, 1992.

[23] D. K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137, 1993.

[24] D. K. Gode and S. Sunder. What makes markets allocatively efficient? *Quarterly Journal of Economics*, 112(2):603–630, 1997.

[25] D. E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison Wesley, 1989.

[26] D. M. Grether. Individual behavior and market performance. *American Journal of Agricultural Economics*, 76(5):1079–1083, 1994.

[27] W. Hardle and A. Kirman. Nonclassical demand: a model-free examination of price-quantity relations in the Marseilles fish market. *Journal of Econometrics*, 67(1):227–257, 1995.

[28] C. A. Holt. Industrial organization: A survey of laboratory research. In J. H. Kagel and A. E. Roth, editors, *The Handbook of Experimental Economics*, pages 349–443. Princeton University Press, Princeton, NJ, 1995.

[29] K. Jamal and S. Sunder. Bayesian equilibrium in double auctions populated by biased heuristic traders. *Journal of Economic Behavior and Organization*, 31(2):273–291, 1996.

[30] J. H. Kagel. Auctions: A survey of experimental research. In J. H. Kagel and A. E. Roth, editors, *The Handbook of Experimental Economics*, pages 501–585. Princeton University Press, Princeton, NJ, 1995.

[31] J. H. Kagel and A. E. Roth, editors. *The Handbook of Experimental Economics*. Princeton University Press, Princeton, NJ, 1995.

[32] J. H. Kagel and W. Vogt. Buyer's bid double auctions: Preliminary experimental results. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories, and Evidence*, pages 285–306. Addison-Wesley, New York, 1992.

[33] A. Leijonhufvud. Towards a not-too-rational macroeconomics. *Southern Economic Journal*, 60(1):1–13, 1993.

[34] G. Y. Luo. Evolution and market competition. *Journal of Economic Theory*, 67(1):223–250, 1995.

[35] G. Mayley. Guiding or hiding: Explorations into the effects of learning on the rate of evolution. In P. Husbands and I. Harvey, editors, *Proceedings of the Fourth European Conference on Artificial Life*, pages 135–144, Cambridge MA, 1997. MIT Press Bradford Books.

[36] R. P. McAfee and J. McMillan. Competition and game-theory. *Journal of Marketing Research*, 33(3):263–267, 1996.

[37] G. F. Miller, P. M. Todd, and S. U. Hegde. Designing neural networks using genetic algorithms. In J. D. Schaffer, editor, *Proceedings of the Third International Conference on Genetic Algorithms*, pages 379–384, San Mateo, California, 1989. Morgan Kaufmann.

[38] A. E. Roth. Introduction to experimental economics. In J. H. Kagel and A. E. Roth, editors, *The Handbook of Experimental Economics*, pages 3–109. Princeton University Press, Princeton, NJ, 1995.

[39] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning internal representations by error propagation. In D. E. Rumelhart and J. L. McClelland, editors, *Parallel Distributed Processing, Volume 1: Foundations*, pages 318–362. MIT Press Bradford Books, Cambridge MA, 1986.

[40] J. Rust, J. Miller, and R. Palmer. Behavior of trading automata in a computerized double auction market. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories, and Evidence*, pages 155–198. Addison-Wesley, New York, 1992.

[41] H. A. Simon. *Sciences of the Artificial*. MIT Press, Cambridge MA, third edition, 1996.

[42] V. L. Smith. An experimental study of competitive market behavior. *Journal of Political Economy*, 70:111–137, 1962.

[43] V. L. Smith. *Papers in Experimental Economics*. Cambridge University Press, Cambridge, 1992.

[44] D. O. Stahl and P. W. Wilson. Experimental evidence on players model of other players. *Journal of Economic Behavior and Organization*, 25(3):309–327, 1994.

[45] K. Steiglitz, M. Honig, and L. Cohen. A computational market model based on individual action. In S. H. Clearwater, editor, *Market-Based Control: A Paradigm for Distributed Resource Allocation*, pages 1–27. World Scientific, Singapore, 1996.

[46] S. Sunder. Experimental asset markets: A survey. In J. H. Kagel and A. E. Roth, editors, *The Handbook of Experimental Economics*, pages 445–500. Princeton University Press, Princeton, NJ, 1995.

[47] M. V. Van Boening and N. T. Wilcox. Avoidable cost: Ride a double-auction roller-coaster. *American Economic Review*, 86(3):461–477, 1996.

[48] G. van Montfort. Economic agents for controlling complex systems. Master's thesis, Delft University of Technology Faculty of Technical Mathematics and Informatics, September 1997. Unpublished.

[49] M. van Tol. Computational agents in a double auction environment: Design, implementation and experimental evidence from an approach with variable information distributions. Master's thesis, Delft University of Technology Faculty of Technical Mathematics and Informatics, July 1998. Unpublished.