# Visualising Search-Spaces for Evolved Hybrid Auction Mechanisms

Dave Cliff
Information Infrastructure Laboratory
HP Laboratories Bristol
HPL-2002-291
October 30th , 2002*

E-mail: dave_cliff@hp.com

automated mechanism design, auctions, marketplaces, ZIP traders GA, genetic algorithm, BICAS

A sequence of previous papers has demonstrated that a genetic algorithm (GA) can be used to automatically discover new optimal auction mechanisms for automated electronic marketplaces populated by software-agent traders. Significantly, the new auction mechanisms are often unlike traditional mechanisms designed by humans for human traders; rather, they are peculiar hybrid mixtures of established styles of mechanism. Qualitatively similar results (i.e., non-standard hybrid mechanism designs being evolved) have been demonstrated for Cliff's ZIP trader algorithm and also for Gode & Sunder's ZI-C traders, provoking the possibility that such hybrid markets may be optimal for *any* marketplace populated entirely by artificial trader-agents. The financial implications of this work could potentially be measured in billions of dollars. In an attempt to elucidate *why* these evolved hybrid markets outperform traditional human-designed mechanisms, this paper presents results from thousands of repetitions of the GA experiments. These data allow 2D projections of the 10-dimensional real-space fitness landscape to be made, which *inter alia* illustrate a surprisingly high sensitivity in the relationship between the fitness evaluation function and the resulting landscape.

* Internal Accession Date Only
Approved for External Publication

# Visualising Search-Spaces for Evolved Hybrid Auction Mechanisms

Dave Cliff
Hewlett-Packard Laboratories, Filton Road, Bristol BS34 8QZ, England, U.K.
dave_cliff@hp.com

**Abstract: A sequence of previous papers has demonstrated that a genetic algorithm (GA) can be used to automatically discover new optimal auction mechanisms for automated electronic marketplaces populated by software-agent traders. Significantly, the new auction mechanisms are often unlike traditional mechanisms designed by humans for human traders; rather, they are peculiar hybrid mixtures of established styles of mechanism. Qualitatively similar results (i.e., non-standard hybrid mechanism designs being evolved) have been demonstrated for Cliff's ZIP trader algorithm and also for Gode & Sunder's ZI-C traders, provoking the possibility that such hybrid markets may be optimal for *any* marketplace populated entirely by artificial trader-agents. The financial implications of this work could potentially be measured in billions of dollars. In an attempt to elucidate *why* these evolved hybrid markets outperform traditional human-designed mechanisms, this paper presents results from thousands of repetitions of the GA experiments. These data allow 2D projections of the 10-dimensional real-space fitness landscape to be made, which *inter alia* illustrate a surprisingly high sensitivity in the relationship between the fitness evaluation function and the resulting landscape.**

*To be presented in abridged form at the "Beyond Fitness: Visualising Evolution" workshop, 8th International Conf. on the Simulation and Synthesis of Living Systems; Sydney, Australia, December 2002.*

## I. INTRODUCTION

ZIP (Zero-Intelligence-Plus) artificial trading agents, introduced in 1997 [1], are software agents (or "robots") that use simple machine learning techniques to adapt to operating as buyers or sellers in open-outcry auction-market environments similar to those used in the experimental economics work of Smith (e.g. [2]). Although initially developed purely to address deficiencies in Gode & Sunder's ZI-C traders [3], recent experimental work by Das *et al.* at IBM [4] has shown that ZIP traders (unlike ZI-Cs) consistently out-perform human traders in human-against-robot auction marketplaces.

The operation of ZIP traders has been successfully demonstrated in experimental versions of continuous double auction (CDA) markets similar to those found in the international markets for commodities, equities, capital, and derivatives; and in posted-offer auction markets similar to those seen in domestic high-street retail outlets [1,2]. In any such market, there are a number of parameters that govern the adaptation and trading processes of the ZIP traders. In the original formulation [1], the values of these parameters were set by hand, using "educated guesses". However, at CIFEr'98, the first results were presented from using a standard genetic algorithm (GA) to automatically optimise these parameter values [5], thereby eliminating the need for skilled human input in deciding the values of the parameters; more details of these GA results were subsequently given in [6].

In all previous work using artificial traders, ZIP or otherwise, the market mechanism (i.e., the type of auction that the traders are interacting within) had been fixed in advance. Well-known market mechanisms from human economic affairs include: the *English Auction* (where sellers stay silent and buyers quote increasing bid-prices; also known as the *ascending-bid mechanism*), the *Dutch Flower Auction* (where buyers stay silent and sellers quote decreasing offer-prices; also known as the *descending-offer mechanism*); the *Vickery* or *second-price sealed-bid* auction (where sealed bids are submitted by buyers, and the highest bidder is allowed to buy, but at the price of the *second-highest* bid -- this curious mechanism encourages honesty and is robust to attack by dishonest means); and the CDA (where sellers announce decreasing offer prices while *simultaneously and asynchronously* the buyers announce increasing bid prices, with the sellers being free to accept any buyer's bid at any time and the buyers being free to accept any seller's offer at any time).

At CIFEr'02, Cliff [7] presented the first results from experiments where a GA optimised not only the parameter values for the trading agents, but also the style of market mechanism in which the traders operate. To do this, a space of possible market mechanisms was created for evolutionary exploration. The space included the CDA and also one-sided auctions similar (but not actually identical to) the English Auction (EA) and the Dutch Flower Auction (DFA); and significantly this space is *continuously variable*, allowing for any of an *infinite* number of peculiar hybrids of these auction types to be evolved, which have no known correlate in naturally occurring market mechanisms. While there was nothing to prevent the GA from settling on solutions that correspond to the known CDA auction type or the EA-like and DFA-like one-sided mechanisms, Cliff [7,8] repeatedly found that the GA settles on hybrid solutions and that these hybrids lead to the most desirable market dynamics. Although the hybrid market mechanisms could easily be implemented in online electronic marketplaces, they have not been designed by humans: rather they are the product of evolutionary search through a continuous space of possible auction-types. Thus, the CIFEr'02 paper [7] was the first demonstration that radically new market mechanisms for artificial traders may be designed by automatic means, thereby establishing the new field of *automated market-mechanism design.* Independently, similar work was under development elsewhere, and was published a couple of months later [9]. Cliff's CIFEr'02 results [7] have recently been independently replicated using a different form of GA and of genetic encoding [10].

As all of Cliff's results [7,8] were from marketplaces populated by ZIP traders, an obvious question to ask is to what extent those results were dependent on the use of ZIP traders.

That is: if non-ZIP trader-agents had been used, would similar hybrid auction mechanisms still be found to be optimal by the GA? This question was answered in the affirmative by Cliff, Walia, & Byde [11,12] who demonstrated qualitatively similar results using Gode & Sunder's ZI-C traders. ZI-C traders are essentially parameter-free stochastic agents, and their computational simplicity made it possible to explicitly plot the "fitness landscapes" showing the economic performance of ZI-C marketplaces over the entire space of possible auction mechanisms being considered. Doing the same for ZIP traders is a much more computationally daunting task, as the trader algorithm is more complex (involving a minimum of eight free real-valued parameters) and also has a richer statistical structure (i.e., it is significantly multimodal). In [11], coarse projections of the ZIP fitness landscapes were produced for comparison with the ZI-C landscapes. In this paper, we present the results of 4000 experiments performed to better illustrate these ZIP fitness landscapes.

The 4000 experiments reported here would represent approximately 8000 hours of continuous computation had a single-CPU computer been used (assuming the use of a 1.8Ghz P4 Hewlett-Packard ePC with 512Mb RAM running Windows2000); that is, roughly 11 months of continuous computation. Luckily, a (shared) compute-cluster built from 50 such ePCs was available for this research, reducing the minimum wait to 160(=8000/50) hours, i.e. a little under one week. Yet by the current publicly-stated aims of various computer companies (including HP), a 50-node cluster is very small fry. Super-clusters of "grid" computers or "utility data centers" are currently planned (or under construction) involving thousands or tens of thousands of such nodes. This paper concludes with some speculative discussion of the implications for artificial evolution research of such a several-orders-of-magnitude increase in available compute resources.

Before that, Section II gives an overview of the background experimental methods and results as published by Cliff [5,6,7,8]; and Section III reviews the visualisations of ZIP-trader fitness landscapes previously published in [11]. Much of Sections II and III is verbatim repetition from our previous papers, but such background is necessary for completeness and intelligibility. In Section IV we then show new visualisations which illustrate how the fitness landscape alters as a key parameter of the evaluation function is gradually changed.

Note that, throughout this paper, $v=\mathbf{U}[x,y]$ is used to denote a random real value $v$ generated from a uniform distribution over the range $[x,y]$.

## II. BACKGROUND

### *A. Zero-Intelligence Plus (ZIP) Traders*

ZIP traders are described fully in [1], which includes sample source-code in the *C* programming language. For the purposes of this paper a high-level description of the key parameters is sufficient. Each ZIP trader $i$ is given a private (secret) limit-price, $\lambda_i$, which for a seller is the price below which it must not sell and for a buyer is the price above which it must not buy. If a ZIP trader completes a transaction at its $\lambda_i$ price then it generates zero utility ("profit" for the sellers or "saving" for the buyers). For this reason, each ZIP trader $i$ maintains a time-varying margin $\mu_i(t)$ and generates quote-prices $p_i(t)$ at time $t$ according to $p_i(t)=\lambda_i(1+\mu_i(t))$ for sellers and $p_i(t)=\lambda_i(1-\mu_i(t))$ for buyers. The "aim" of traders is to maximise their utility over all trades, where utility is the difference between the accepted quote-price and the trader's $\lambda_i$ value. Trader $i$ is given an initial value $\mu_i(0)$ (i.e., $\mu_i(t)$ for $t=0$) which is subsequently adapted over time using a simple machine learning technique known as *the Widrow-Hoff rule* which is also used in back-propagation neural networks. This rule has a "learning rate" parameter $\beta_i$ that governs the speed of convergence between trader $i$'s quoted price $p_i(t)$ and the trader's idealised "target" price $\tau_i(t)$. When calculating $\tau_i(t)$, traders introduce a small random absolute perturbation generated from $\mathbf{U}[0,c_a]$, and also a small random relative perturbation coefficient generated from $\mathbf{U}[1-c_r,1]$ (when a trader is reducing its $p_i(t)$) or $\mathbf{U}[1,1+c_r]$ (when increasing $p_i(t)$) where $c_a$ and $c_r$ are global system constants. To smooth over noise in the learning system, there is an additional "momentum" parameter $\gamma_i$ for each trader (such momentum terms are also commonly used in back-propagation neural networks).

Thus, adaptation in each ZIP trader $i$ has the following parameters: initial margin $\mu_i(0)$; learning rate $\beta_i$; and momentum term $\gamma_i$. In an entire market populated by ZIP traders, these three parameters are assigned to each trader from uniform random distributions each of which is defined via "min" and "delta" values in the following fashion: $\mu_i(0)=\mathbf{U}(\mu_{min}, \mu_{min}+\mu_\Delta)$; $\beta_i=\mathbf{U}(\beta_{min}, \beta_{min}+\beta_\Delta)$; and $\gamma_i=\mathbf{U}(\gamma_{min},\gamma_{min}+\gamma_\Delta)$; $\forall i$.

Hence, to initialise an entire ZIP-trader market it is necessary to specify values for the six market-initialisation parameters $\mu_{min}$, $\mu_\Delta$, $\beta_{min}$, $\beta_\Delta$, $\gamma_{min}$, and $\gamma_\Delta$; and also for the two system constants $c_a$ and $c_r$. And so it can be seen that any set of initialisation parameters for a ZIP-trader market exists within an eight-dimensional real space, conventionally denoted by $\mathbf{R}^8$. Vectors in this 8-space can be considered as genotypes, and from an initial population of such genotypes it is possible to allow a GA to find new genotypes that best satisfy an appropriate evaluation function. This is exactly the process that was introduced at CIFEr'98 [5,6], as described in Section II.C below. Before that, we discuss the issue of simulating the passage of time.

When monitoring events in a real auction, as more precision is used to record the time of events, so the likelihood of any two events occurring at exactly the same time is diminished. For example, if two quotes made at five minutes past nine are both recorded as occurring at 09:05, then they appear in the record as simultaneous; but a more accurate clock would have been able to reveal that the first was made at 09:05:01.64 and the second at 09:05:01.98. Even if two events occur absolutely at the same time, very often some

random process (e.g. what direction the auctioneer is looking in) acts to break the simultaneity.

Thus, we may simulate real marketplaces (and implement electronic marketplaces) using techniques where each significant event always occurs at a unique time. We may choose to represent these by real high-precision times, or we may abstract away from precise time-keeping by dividing time into discrete (possibly irregular) *slices*, numbered sequentially, where one significant event is known to occur in each slice. Such a time-slicing approach was used in previous work [1,5,6,7,8,11]. In each time-slice, the atomic "significant event" is one quote being issued by one trader and the other traders then responding either by ignoring the quote or by one of the traders accepting the quote. (NB in [4] a continuous-time formulation of the ZIP-trader algorithm was used).

In the markets described here and in [1,5,6,7,8,11], on each time-slice a ZIP trader $i$ is chosen at random from those currently able to quote (i.e. those who hold appropriate stock or currency), and trader $i$'s quote price $p_i(t)$ then becomes the "current quote" $q(t)$ for time $t$. Next, all traders $j$ on the contraside (i.e. all buyers $j$ if $i$ is a seller, or all sellers $j$ if $i$ is a buyer) compare $q(t)$ to their own current quote price $p_j(t)$ and if the quotes cross (i.e. if $p_j(t)<=q(t)$ for sellers, or if $p_j(t)>=q(t)$ for buyers) then the trader $j$ is able to accept the quote. If more than one trader is able to accept, one is chosen equiprobably at random to make the transaction. If no traders are able to accept, the quote is regarded as "ignored". Once the trade is either accepted or ignored, the traders update their $\mu(t)$ values using the learning algorithm outlined above, and the current time-slice ends. This process repeats for each time-slice in a trading period, with occasional injections of fresh currency and stock, or redistribution of $\lambda_i$ limit prices, until either a maximum number of transactions have occurred, or until either no seller or no buyer is able to quote, or until a maximum number of time-slices have passed since the last accepted quote (i.e., a until a protracted sequence of successive ignored quotes occurs). This protracted-sequence-of-ignored-quotes termination criterion is of relevance to the discussion in Section IV of this paper, so we should note here that in the implementation used for this work (as in the version published in [1]) the number of ignored quotes necessary to terminate trading is set by the system constant MAX_FAILS, which has been set to 100 in the work reported here, as in the ZIP work reported previously [1,5,6,7,8,11].

### *B. Space of Possible Auctions*

Now consider the case where we implement a ZIP-trader continuous double auction (CDA) market. In any one time-slice in a CDA either a buyer or a seller may quote, and in the definition of a CDA a quote is equally likely from each side. One way of implementing a CDA is, at the start of each time-slice, to generate a random binary variable to determine whether the quote will come from a buyer or a seller, and then to randomly choose one individual as the quoter from whichever side the binary value points to. Here, as in previous ZIP work [1,5,6] the random binary variable is always independently and identically distributed over all time-slices.

Let $Q=b$ denote the event that a buyer quotes on any one time-slice and let $Q=s$ denote the event that a seller quotes, then for the CDA we can write $Pr(Q=s)=0.5$ and note that because $Pr(Q=b)=1.0-Pr(Q=s)$ it is only necessary to specify $Pr(Q=s)$, which we will abbreviate to $Q_s$ hereafter. Note additionally that in an EA we have $Q_s=0.0$, and in the DFA we have $Q_s=1.0$. Thus, there are at least three values of $Q_s$ (i.e. $0.0$, $0.5$, and $1.0$) that correspond to three types of auction familiar from centuries of human economic affairs.

However, although the ZIP-trader case of $Q_s=0.5$ is indeed a good approximation to the CDA, the fact that any ZIP trader $j$ will accept a quote whenever $q(t)$ and $p_j(t)$ cross means that the one-sided extreme cases $Q_s=0.0$ and $Q_s=1.0$ are not exact analogues of the EA and DFA. Nevertheless, consider the implications of considering values of $Q_s$ of 0.0, 0.5, and 1.0 not as three distinct market mechanisms, but rather as three points on a *continuum*. How do we interpret, for example, $Q_s=0.1$? Certainly there is a straightforward implementation: on the average, for every nine quotes by buyers, there will be one quote from a seller. Yet the history of human economic affairs offers no examples (as far as we are aware) of such markets: why would anyone suggest such a bizarre way of operating, and who would go to the trouble of arbitrating (i.e., acting as an auctioneer for) such a mechanism? Nevertheless, there is no *a priori* reason to argue that the three known points on this $Q_s$ continuum are the only loci of useful auction types. Maybe there are circumstances in which values such as $Q_s=0.1$ are preferred. Given the infinite nature of a real continuum, it seems appealing to use an automatic exploration process, such as a GA, to identify useful $Q_s$ values.

Thus, a ninth dimension was added to the search space, and the genotype in the GA is now the eight real values for ZIP-trader initialisation, plus a real value for $Q_s$, so the GA is searching for points in $\mathbf{R}^9$ that give the best market dynamics.

### *C. The Genetic Algorithm*

A simple genetic algorithm was used. As with each experiment reported in [5,6,7,8,11] a population of size 30 was used, and evolution was allowed to progress for some number of generations $n_g$. In each generation, all individuals were evaluated and assigned a "fitness" value (reflecting how good that genotype's market dynamics were); and the next generation's population was then generated via mutation and crossover on parents identified using rank-based tournament selection. Elitism (where an unaltered copy of the fittest individual from generation $g$ is inserted into the population of $g+1$) was also used.

The genome of each individual was simply a vector of nine real values. In each experiment, the initial random population was created by generating random values from $\mathbf{U}[0,1]$ for

each locus on each individual's genotype. Crossover points were between the real values, and crossover was governed by a Poisson random process with an average of between one and two crosses per reproduction event. Mutation was implemented by adding random values from $\boldsymbol{U}[-m(g),+m(g)]$ where $m(g)$ is the mutation limit at generation $g$ (starting the count at $g=0$). Mutation was applied to each locus in each genotype on each individual generated from a reproduction event, but the mutation limit $m(g)$ was gradually reduced via an exponential-decay annealing function of the form: $log_{10}(m(g))=-(log_{10}(m_s)-(g/(N_g-1))log_{10}(m_s/m_e))$ where $N_g$ is the maximum number of generations and $m_s$ is the "start" mutation limit (i.e., for $m(0)$) and $m_e$ is the "end" mutation limit (i.e., for $m(n_g-1)$). In all the experiments reported here and in earlier papers [7,8], $N_g=10^3$, $m_s=0.05$, and $m_e=0.0005$.

If ever mutation caused the value at a locus to fall outside $[0.0,1.0]$ it was simply clipped to stay within that range. This clip-to-fit approach to dealing with out-of-range mutations biases evolution toward extreme values (i.e. the upper and lower bounds of the clipping), and so $Q_s$ values of 0.0 or 1.0 are, if anything, more likely to be evolved. Initial and mutated genome values of $\mu_\Delta$, $\beta_\Delta$, and $\gamma_\Delta$ were also clipped to satisfy $(\mu_{min}+\mu_\Delta)<=1.0$, $(\beta_{min}+\beta_\Delta)<1.0$, & $(\gamma_{min}+\gamma_\Delta)<1.0$.

The fitness of genotypes was evaluated using the same methods as described in [5,6,7,8,11]. One *trial* of a particular genome was performed by initialising a ZIP-trader market from the genome, and then allowing the traders to operate within the market for a fixed number of trading periods, with allocations of stock and currency being replenished between trading periods. Each trading period was ended in the manner described at the end of Section II.A.

During each trading period, Smith's $\alpha$ measure [2] of deviation of transaction prices from the theoretical market equilibrium price was monitored, and a front-weighted average was calculated across the trading periods in the trial. As the outcome of any one such trial is influenced by stochasticity in the system, the final fitness value for an individual was calculated as the arithmetic mean of 100 such trials. Note that as minimal deviation of transaction prices from the theoretical equilibrium price is desirable, lower scores are better: the intention here is to *minimise* the fitness value.

In [7] the number of generations $n_g$ for each experiment was set to equal $N_g$ (i.e., 1000), but all the significant evolutionary activity was found to occur in the first 500 generations; hence in subsequent work [8,11] $n_g$=500 was used. Thus, in any one experiment, there are 30 individuals evaluated over 500 generations where each evaluation involves calculating the mean of 100 trials, so a total of 1.5 million market trials would be executed in any one GA experiment. Nevertheless, the progress of each GA experiment is itself affected by stochasticity (e.g. the GA may become trapped on local optima) and so to generate reliable results each experiment was repeated 50 times, requiring a total of 75 million market trials. On a current single-CPU PC, 50 repetitions of the single-schedule experiments from [7] take around four days to complete. As shown by Walia [12], much of this consumption of compute time is due to our choice of a computationally expensive (but statistically rigorous) random number generator function.

*D. Previous Results*

In the CIFEr'02 paper [7], three differing market supply and demand schedules were used, shown here in Figures 1, 2, and 3, and referred to hereafter as markets M1, M2, and M3 respectively.

Each of Figures 1 to 3 shows a supply and demand schedule for a marketplace with 11 buyers and 11 sellers, each empowered to buy/sell one unit of commodity, and all three are similar (or identical) to the schedules used by Smith [2]. Figure 4 shows results from 50 repetitions of an experiment where the GA explores the $\mathbf{R}^9$ subspace in an attempt to optimise the ZIP-trader market parameters for operating in M1: for each experiment, the fitness of the best (elite) member of the population is recorded. The results are clearly tri-modal.

Of the 50 repetitions, in five the elite ends up on fitness minima of about 3.2, while the other two elite fitness modes are on less-good minima of around 4.0 and 4.75. For comparison, Figure 5 shows the results of 50 repeats of the same experiment, where the value of $Q_s$ was *not* evolved, being instead clamped at 0.5: i.e. the CDA value. The CDA mechanism is often applauded as an auction mechanism in which equilibration is rapid and stable, so we could expect the best fitness from using this market type. With the fixed CDA auction style, an average elite fitness of around 4.5 is settled on by the majority of experiments (48 repetitions) while a small minority (2 repetitions) settle on a less good mode of around 5.1. Clearly then, the evolved-mechanism results are better than the fixed-mechanism CDA results; that is, when the GA is allowed so find its own value of $Q_s$ rather than have the CDA $Q_s$ value of 0.5 imposed on it, it finds fitter solutions – solutions with less deviation of transaction prices from the equilibrium price. As it happens, the $Q_s$ value found in the best elite mode for the evolving-mechanism M1 experiments is zero [7], and for M2 the best $Q_s$ was also zero [7]. But, surprisingly and significantly, for M3 the best $Q_s$ was neither zero, nor 0.5, nor 1.0 – i.e. none of the $Q_s$ values corresponding to traditional human-designed auction mechanisms; rather, the best $Q_s$ for M3 was found to be around 0.16 [7].

All of the results in the CIFEr'02 paper came from experiments in which the same static supply/demand schedule was used for the duration of each evaluation of every genotype. This is a somewhat unrealistic simplification, for two reasons. First, a primary reason why auction mechanisms such as the CDA are of interest is their ability to adapt to changes in the market's supply and demand curves. Second, it is likely that the GA exploited this regularity and *over-fitted* the ZIP-trader parameters to the particular market schedules used (e.g. a genome that does well in M1 may perhaps perform

poorly in M2). Thus, in a subsequent paper [8], similar experiments were run but in these new studies the evaluation of a genotype involved six trading periods on one market schedule, followed by a shock-change to another schedule, and then another six trading periods on the new schedule; with the fitness of the genotype being calculated over the entire twelve periods of trading.

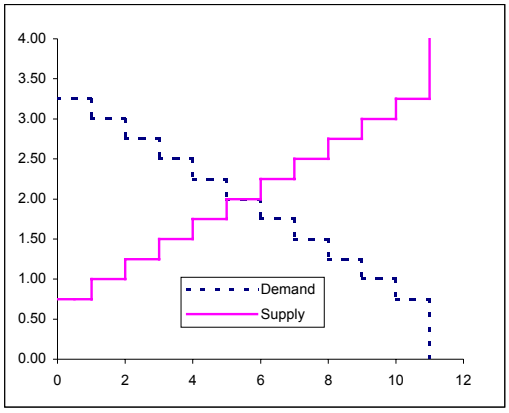


Figure 1: Supply and demand schedules for market M1. Vertical axis is Price; horizontal axis is Quantity.

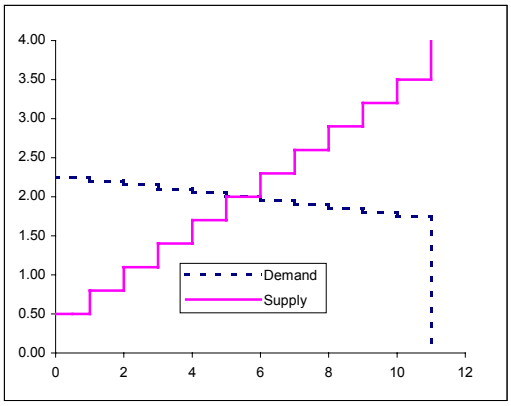


Figure 2: Supply and demand schedules for market M2.

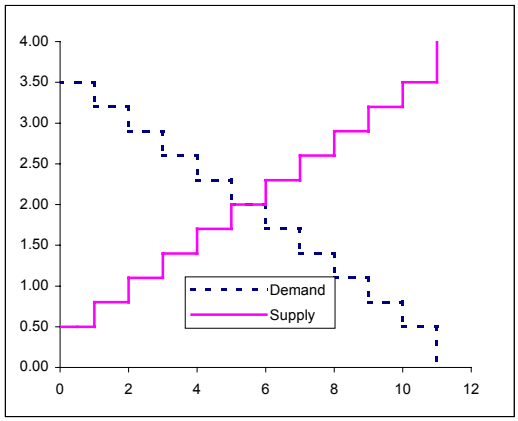


Figure 3: Supply and demand schedules for market M3.

Hence in these experiments the genotypes had to optimise not only the ZIP-trader's *ab initio* adaptation to the first schedule but also their re-adaptation to the new schedule introduced half-way through the evaluation process. Two sets of experiments were performed: one in which the ZIP traders operated in M1 for six periods followed by a shock-change to M2 for the final six periods (referred to as the M1M2 experiments); and another in which the order was M2 followed by M1 (referred to as M2M1). It was demonstrated [8] that the order was significant: the M1M2 results differed significantly from the M2M1 results. Although in the single-schedule experiments both M1 and M2 were found to have optima at $Q_s$=0, when the two schedules were both used in one trial then non-zero values of $Q_s$ evolved: for M1M2 the best-mode value was a "hybrid" of around 0.25; while for M2M1 the best value was 0.45, which did not yield statistically significant differences in performance from the CDA value of 0.5.

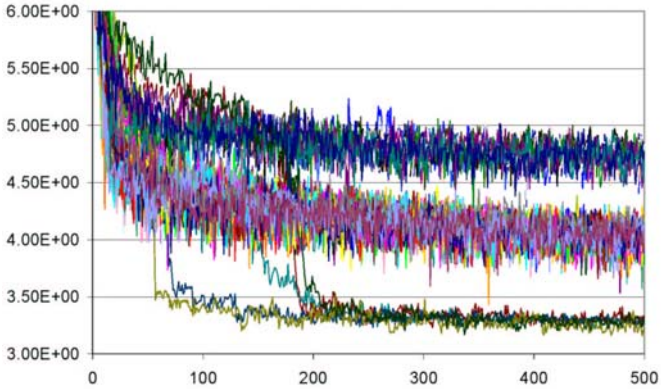


Figure 4: Elite fitness values from 50 repetitions of a 500-generation evolving-mechanism (EM) experiment operating with M1. Lower values are better solutions (less deviation from equilibrium). Results are trimodal, with five of the repetitions (10% ) settling to values around 3.2.

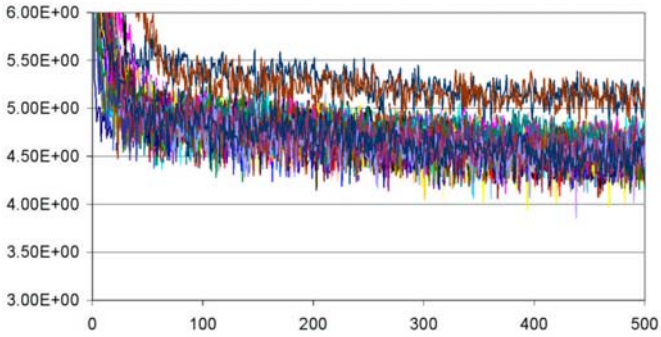


Figure 5: Elite fitness values from 50 repetitions of a 500-generation experiment operating with M1, but with a fixed-mechanism (FM) CDA of $Q_s$=*0.5*: bimodal results, with 96% of the repetitions settling to fitness values around 4.5 and the remaining 4% at around 5.2.

## III. ZIP-TRADER FITNESS LANDSCAPES

### *A. Methods and Results*

As was stated in the previous section, the genotypes in the ZIP-trader experiments are within $\mathbf{R}^9$ (strictly, they are all within the real unit hypercube $[0.0,1.0]^9$). For any such genotype, one evaluation (e.g. taking the mean score from 100 trials, as used here) will give a fitness score for that genotype; and so it is possible in principle to visualise the "fitness landscape" as a surface over the 9-d axes of the genotype space. Visualising such a 10-d object in the two or three dimensions that we humans are familiar with communicating in is manifestly problematic; yet appropriate visualisations can be highly valuable in demonstrating that the results of the GA's evolutionary search are indeed a plausible global optimum. Thus, in this section, we present data (first published in [11]) showing visualisations of the fitness landscapes for all five of the ZIP-trader experiments reported in [7] and [8] (i.e., M1, M2, M3, M1M2, and M2M1), before showing in Section IV the fitness landscapes from comparable experiments where the markets are populated by ZI-C traders.

To understand the visualisation, consider Figure 5. In this set of fifty M1 experiments the value of $Q_s$ was fixed at 0.5 and by generation 500 there are two clear elite-fitness modes: one at approx 4.5 and one at approx 5.2. Of the fifty repetitions, 96% settle to the first mode and 4% settle to the second. This could be represented by a histogram where the horizontal axis represents discretized ("binned") values of the elite-fitness mode, and the vertical axis represents the frequency with which each mode is observed; for the M1 $Q_s$=0.5 data of Figure 5 we would see two distinct peaks in the histogram: a big one around 4.5 and a smaller one around 5.2.

Now to visualise the entire fitness landscape for ZIP traders in M1, run more fixed-mechanism experiments but for each set of 50 repetitions hold the value of $Q_s$ fixed at some value while all the other 8 ZIP parameters on the genome are optimised by the GA. These data allow us to plot a 3-d projection of the 10-d fitness landscape: in our projection, one horizontal dimension is the fixed value of $Q_s$; another is the elite-fitness mode-value; and the vertical axis shows the frequency with which the different mode values are reached for each of the $Q_s$ values. Figure 6 shows a perspective projection of such a 3-d histogram, calculated for ZIP traders in M1. An alternative view of the same data is presented in Figure 7, i.e. as a contour plot with a logarithmic compression function applied to the frequency values. For comparison, Figures 8, 9, 10, and 11 show such contour plots of the fitness landscape for ZIP-trader experiments with markets M2, M3, M1M2, and M2M1 respectively.

### *B. Discussion*

All of the contour plots of the fitness landscapes show good agreement with the previous experimental results, in the sense that the minima (i.e. left-most data points) on the contour plots are in good agreement with the values discovered by the evolving-mechanism GA experiments reported on in [7,8] and summarised in Section II.D: markets M1 and M2 both have minima at $Q_s$=0.0; M3 has a minimum between $Q_s$=0.125 and $Q_s$=0.25; M1M2 has a clear minimum at $Q_s$=0.25; and M2M1 has a minimum at $Q_s$=0.5. It is also worth noting that all the contour plots show some degree of multi-modality for some values of $Q_s$.

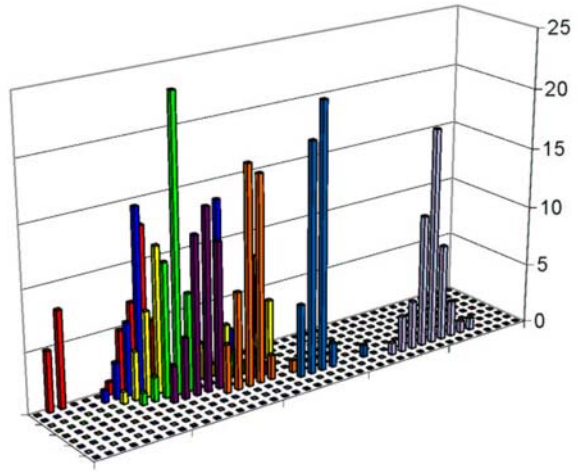


Figure 6: 3-d histogram showing elite fitness-mode frequency data from multiple repetitions of ZIP-trader GA experiments in M1 over a variety of fixed $Q_s$ values. The narrow horizontal axis is $Q_s$ (from 0.0 at the rear to 1.0 at the front, increasing at intervals of 0.125); the long horizontal axis is elite-fitness value from 3.0 at the left to 8.0 at the right, in "bins" of 0.125; the vertical axis shows the frequency with which the GA settles to each elite-fitness (out of 50 repetitions at each fixed $Q_s$ value). Note that the elite-fitness values for $Q_s$=1.0 are so poor (i.e., so high) that their histogram data lie off the scale to the right.

Thus, the evolving-mechanism results from [7,8] are supported by this brute-force exploration of the fixed-mechanism fitness landscapes for each market schedule: in each case, the evolved value of $Q_s$ is very close to the value identified by empirical examination of the fitness landscapes, and the multi-modality of each fitness landscape justifies the use of multiple repetitions of each experiment in order to identify the true optimal solution.

However, the differences between Figure 7 (M1) and Figure 9 (M3) are intriguing: the two contour plots show some major differences, yet examination of the supply and demand schedules (shown in Figures 1 and 3 respectively) reveals that these differences in the fitness landscapes are caused by only minor changes in the supply and demand schedules. Specifically, in both M1 and M3 the equilibrium price is 2.00 and the equilibrium quantity is 6; both schedules have 11 buyers and 11 sellers; and both schedules have identical differences in limit prices between successive traders. In fact, the *only* difference between M1 and M3 concerns the magnitude of the difference in limit prices between successive traders, which we'll denote by $\Delta\lambda$: in M1, $\Delta\lambda$=0.25; in M3, $\Delta\lambda$=0.30.

That such an apparently minor difference in the market schedules can have a relatively major effect on the resultant fitness landscapes is a phenomenon deserving of more detailed exploration, which is the topic of the rest of this paper.

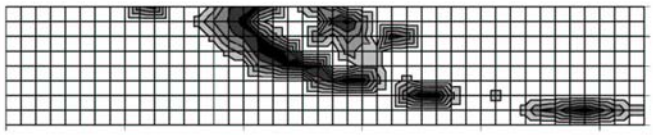

Figure 7: Contour plot of the data shown in Figure 6. Horizontal axis is fitness values from 2.0 at the left to 7.5 at the right (grid-spacing is 0.125); vertical axis is $Q_s$ from 0.0 at the top to 1.0 at the bottom (grid-spacing is 0.125). Darker shading represents higher frequency. Nonzero frequencies for $Q_s$=1.0 are so poor that they lie off the scale to the right.

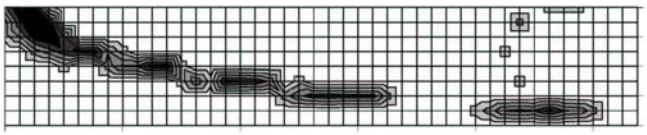

Figure 8: Contour plot of fitness landscape for ZIP traders in M2. Scale as for Fig. 7. Nonzero frequencies for $Q_s$=1.0 lie off the scale to the right.

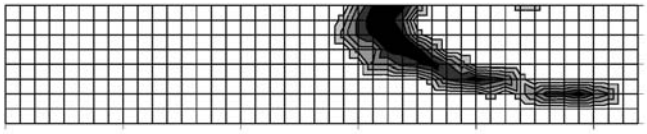

Figure 9: Contour plot of fitness landscape for ZIP traders in M3. Scale as for Figure 7. Fitness values for $Q_s$>=0.875 lie off the scale to the right.

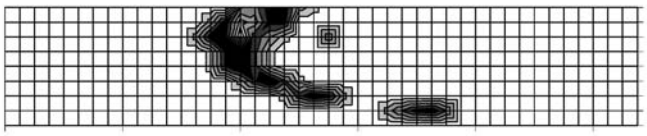

Figure 10: Contour plot of fitness landscape for ZIP traders in M1M2. Scale as for Figure 7. Nonzero frequencies for $Q_s$=1.0 lie off the scale to the right.

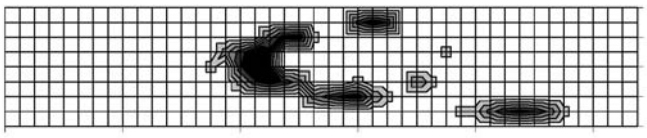

Figure 11: Contour plot of fitness landscape for ZIP traders in M2M1. Scale as for Figure 7. Nonzero frequencies for $Q_s$=0.0 and $Q_s$=1.0 lie off the scale to the right.

Having established the background to our current work, we now proceed with introducing our new results. In the next section we show finer-grained projections of the fitness landscapes that result from using schedules M1 and M3, and also from intermediate schedules where the step-size Δλ has values between 0.25 and 0.30.

## IV. HIGH RESOLUTION ZIP FITNESS LANDSCAPES

### *A. Methods*

The contour plots presented in the previous section are one approach to visualising densities of the elite-fitness modes, but the fact that the data was calculated at fixed discrete $Q_s$ values (of 0.0, 0.125, 0.250, … , 1.00) introduces some unappealing artefacts. For instance, in Figure 8 the "stepped" appearance of the main high-frequency peak ridge from $Q_s$=0.375 to $Q_s$=0.750 is an artefact (the true underlying ridge is probably smooth, but appears jagged because it is sampled at discrete intervals). The apparent jump-discontinuity in that peak-ridge from $Q_s$=0.750 to $Q_s$=0.875 is also artefactual: a finer-grained sampling of $Q_s$ values (e.g. at $Q_s$=0.755, 0.760, 0.765, etc) would have provided data to fill in this "gap".

A second problem with the contour plots is that they show histograms for the elite-fitness data at generation 500, but give no indication of the variability in that data (i.e., how different would the contour plot have looked at generation 499, or 475?). In general, it is more rigorous and more informative to show measures of central tendency (e.g. a mean) and of variability (e.g. a standard deviation).

To avoid these visualisation artefacts, and to show the data in a more rigorous fashion, a new set of visualisation experiments was conducted. In each experiment, a fixed value for $Q_s$ was chosen at random from a uniform distribution, and that value remained constant throughout the experiment. All other experiment details were the same as those described in Section II, except that the final measure of an experiment's outcome is now the mean and standard deviation of the elite fitness over the final 50 generations of the experiment (i.e., generations 450 to 500). The result of any one such experiment can be plotted on a 2D graph showing $Q_s$ on the horizontal axis, with the vertical axis showing mean elite fitness (and error bars indicating the standard deviation on that mean). Performing a large number of repetitions of this experiment (e.g. $n$=500) with a different randomly-generated value of $Q_s$ on each repetition gives a scatter-plot which as a visualization technique is broadly similar to the contour plots shown previously but which avoids their problems. In the remainder of this paper, $f_\mu$ will be used to denote mean elite fitness.

### *B. Results*

$Q_s/f_\mu$ scatter-plots were generated for Δλ=0.25, 0.26, 0.27, 0.28, 0.29, and 0.30. They are illustrated in Figures 12 to 17 respectively. The data for M1 (Δλ=0.25) in Figure 12 shows three clear modes. A dominant mode runs from $f_\mu$=~4.1 at $Q_s$=0.0 through to $f_\mu$=~4.3 at $Q_s$=0.5. For later discussion, we'll call this Mode A. Above and roughly parallel to Mode A there lies a second weaker mode running from $f_\mu$=~4.75 at $Q_s$=0.0 to $f_\mu$=~5.00 at $Q_s$=0.5, which we'll call Mode B. Fi-

nally, a third mode sweeps up and to the left from $f_\mu$=~3.5 at $Q_s$=0.00 up to $f_\mu$=~5.5 at $Q_s$=0.25, beyond which it rapidly fades. We'll call that Mode C.

Examining Figures 13 to 17, we see that each change of 0.01 in λΔ has a noticeable systematic change on Modes A, B, and C. At a macro-level, we see that Modes B and C act as less

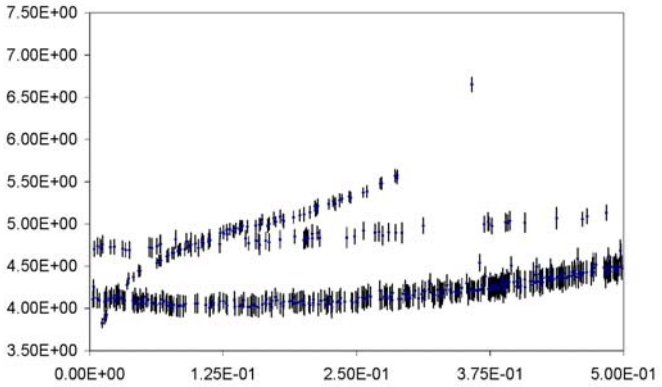


Figure 12: λΔ=0.25 $Q_s/f_\mu$ scatter-plot. Vertical axis is fitness, horizontal axis is fixed $Q_s$ value. Each data-point marks the mean elite fitness $f_\mu$ from the final 50 generations of a 500-generation experiment, with the vertical error bars indicating plus and minus one standard distribution. This plot shows data from 500 repetitions of the experiment with the fixed $Q_s$ value for each experiment being generated from $Q_s$=***U**[0.0,0.5]*.

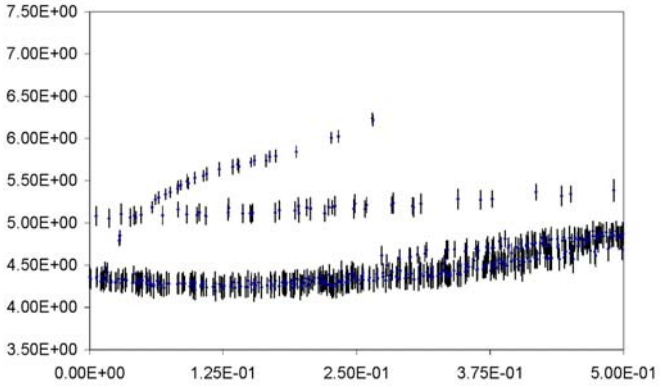


Figure 13: λΔ=0.26 $Q_s/f_\mu$ scatter-plot. Format as for Figure 12.

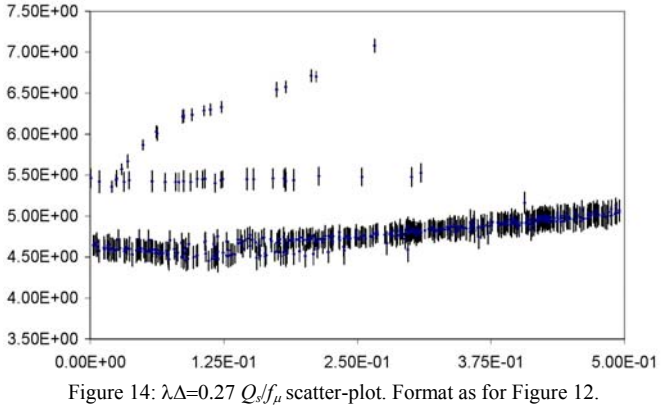


Figure 14: λΔ=0.27 $Q_s/f_\mu$ scatter-plot. Format as for Figure 12.

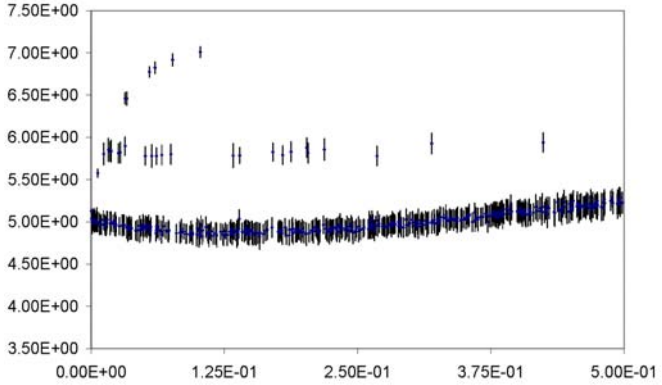


Figure 15: λΔ=0.28 $Q_s/f_\mu$ scatter-plot. Format as for Figure 12.

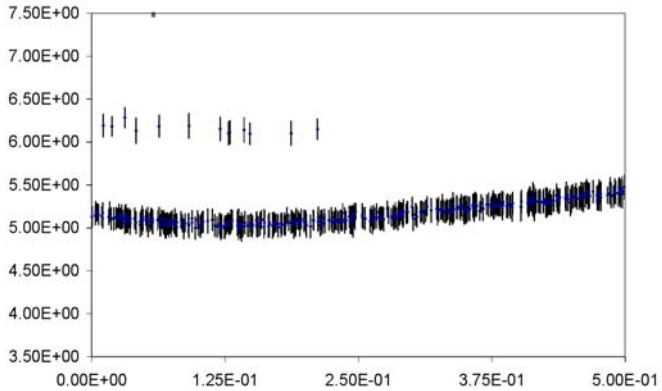


Figure 16: λΔ=0.29 $Q_s/f_\mu$ scatter-plot. Format as for Figure 12.

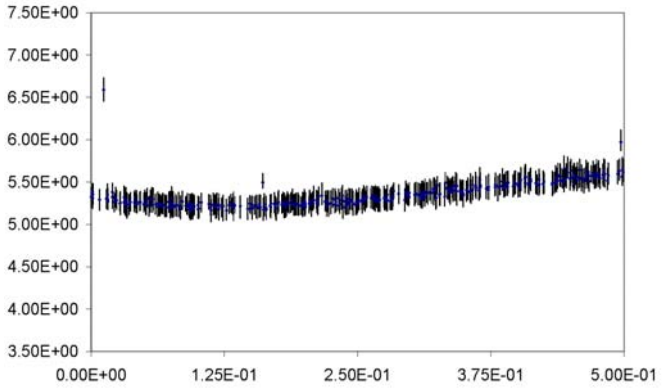


Figure 17: λΔ=0.30 $Q_s/f_\mu$ scatter-plot. Format as for Figure 12.

powerful attractors as Δλ increases, so that by λΔ=0.30 (i.e., M3) Mode A appears to be effectively all that remains. At the micro-level, in the change from λΔ=0.25 to λΔ=0.26, Modes A and B have both moved up slightly, while Mode C has moved upwards more rapidly. Furthermore, for $Q_s$=0.25 to $Q_s$=0.50, a new mode (which we'll call mode D) is appearing slightly above mode A. The distinction between these Modes A and D is seen more clearly when the error bars are removed from the scatter plots. Thus, Figures 18 to 23 show the same data as Figures 12 to 17, but with no error-bars and with the vertical axis scales altered for greater clarity.

Figures 19 and 20 indicate that as λΔ increases, Mode D becomes a stronger attractor; such that by λΔ=0.28 the dominant mode is actually D rather than A.

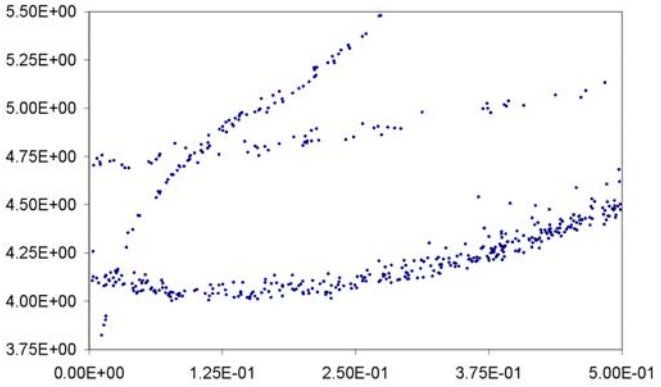


Figure 18: λΔ=0.25 $Q_s/f_\mu$ scatter-plot as shown in Figure 12 but with error-bars removed and vertical-axis scale altered.

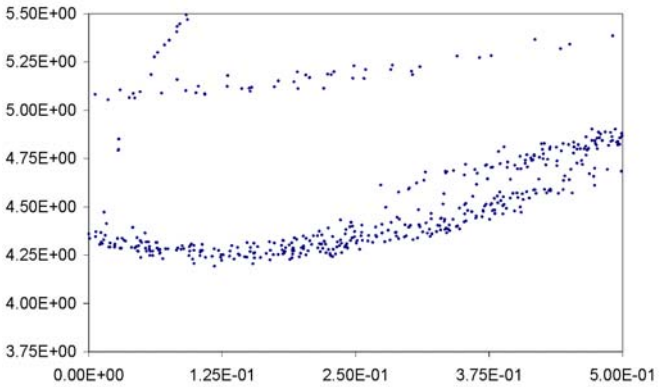


Figure 19: λΔ=0.26 $Q_s/f_\mu$ scatter-plot as shown in Figure 13 but with error-bars removed and vertical-axis scale altered.

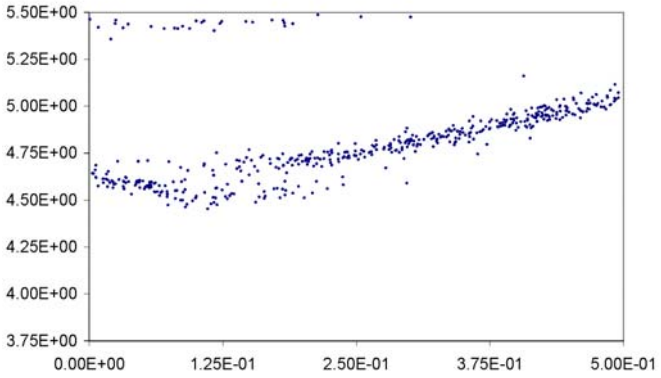


Figure 20: λΔ=0.27 $Q_s/f_\mu$ scatter-plot as shown in Figure 14 but with error-bars removed and vertical-axis scale altered.

To establish that these different $f_\mu$ modes do represent distinct modes in genotype space, we can seek to identify distinct clusters in the $\mathbf{R}^8$ genotype space (recall, in these experiments the 9th genotype dimension, i.e. $Q_s$, was clamped to a random value for each experiment.) Although a clustering algorithm could be used to analyse the genome data, as it happens the distinct $f_\mu$ modes are clearly (and conveniently) strongly correlated with one locus on the genotype, allowing the clusters to be identified by inspection. This locus is the $\mu_{min}$ (base profit margin at time zero) value. Figures 24 to 29 show the $\mu_{min}$ data for λΔ=0.25 to 0.30 respectively.

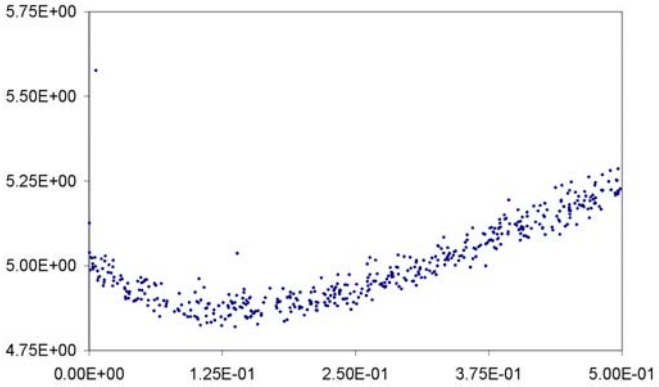


Figure 21: λΔ=0.28 $Q_s/f_\mu$ scatter-plot as shown in Figure 15 but with error-bars removed and vertical-axis scale altered.

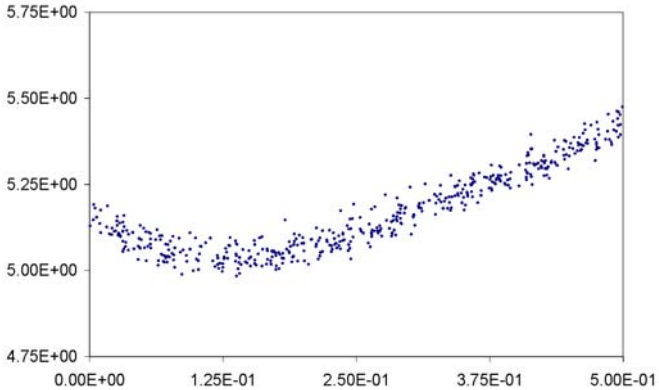


Figure 22: λΔ=0.29 $Q_s/f_\mu$ scatter-plot as shown in Figure 16 but with error-bars removed and vertical-axis scale altered.

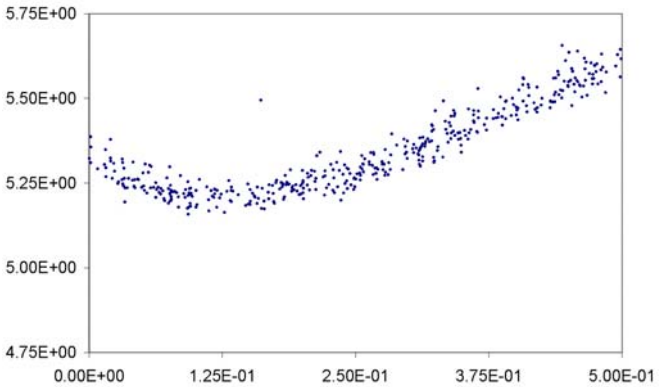


Figure 23: λΔ=0.30 $Q_s/f_\mu$ scatter-plot as shown in Figure 17 but with error-bars removed and vertical-axis scale altered.

Figures 24 to 29 show that genomes in Mode C all have $\mu_{min}$ =~1.0 and that as λΔ increases Mode C weakens, failing to attract any of the elite genotypes once λΔ>0.28. Mode B can

also be seen as the line of genomes with $\mu_{min}$ values spread from ~0.0 to ~0.25 in Figure 24; a line which rises and fades as λΔ increases in Figures 25 to 29. The interplay between Modes A and D as λΔ increases is made more clear when the same co-ordinate data is normalised by the underlying value of λΔ: that is, when the mean elite fitness value and the mean elite $\mu_{min}$ value for each data-point are divided by the λΔ for the experiment. This normalisation accounts for the fact that

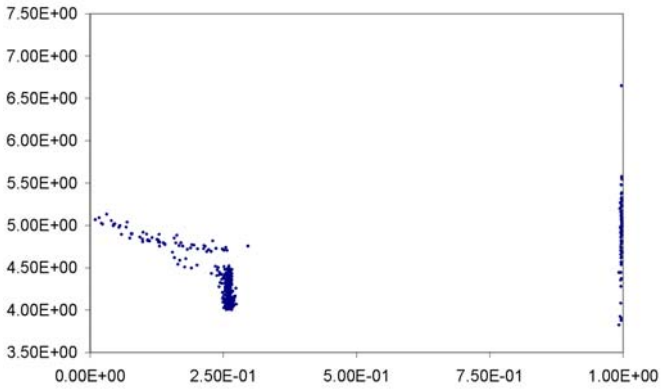


Figure 24: λΔ=0.25 $\mu_{min}/f_{\mu}$ scatter-plot. Horizontal axis is genome $\mu_{min}$ value; vertical axis is fitness. Each data-point shows the mean elite fitness over generations 450 to 500 of an experiment, plotted against the mean elite genome $\mu_{min}$ value over those 50 generations. Data from the 500 repetitions shown in Figures 12 and 18.

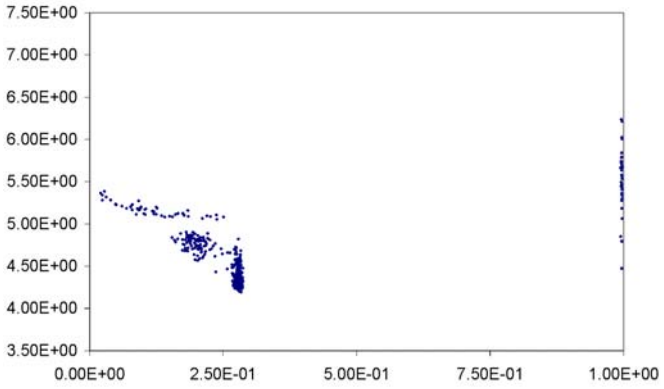


Figure 25: λΔ=0.26 $\mu_{min}/f_{\mu}$ scatter-plot; format as for Figure 24; data from the 500 repetitions illustrated in Figures 13 and 19.

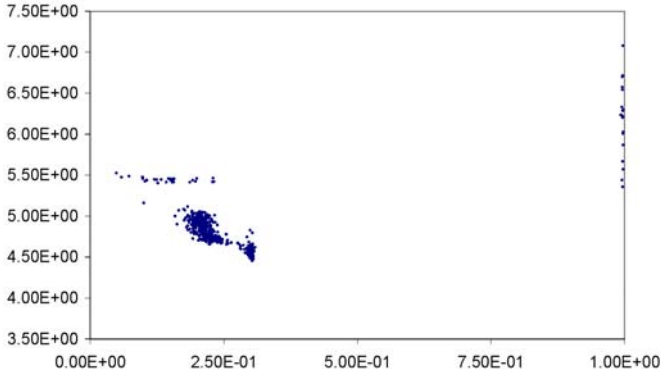


Figure 26: λΔ=0.27 $\mu_{min}/f_{\mu}$ scatter-plot; format as for Figure 24; data from the 500 repetitions illustrated in Figures 14 and 20.

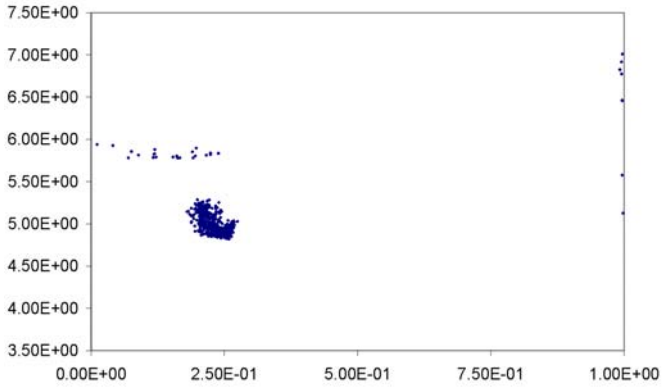


Figure 27: λΔ=0.28 $\mu_{min}/f_{\mu}$ scatter-plot; format as for Figure 24; data from the 500 repetitions illustrated in Figures 15 and 21.

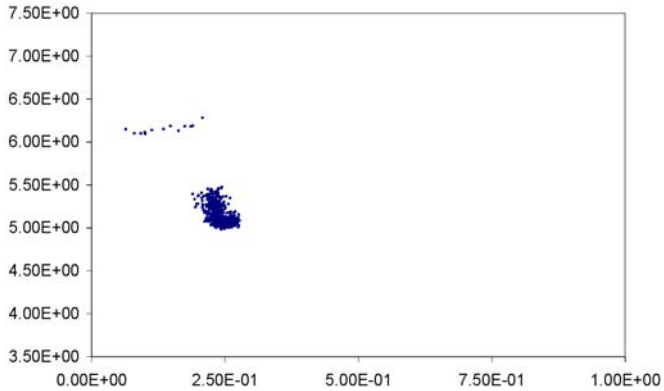


Figure 28: λΔ=0.29 $\mu_{min}/f_{\mu}$ scatter-plot; format as for Figure 24; data from the 500 repetitions illustrated in Figures 16 and 22.

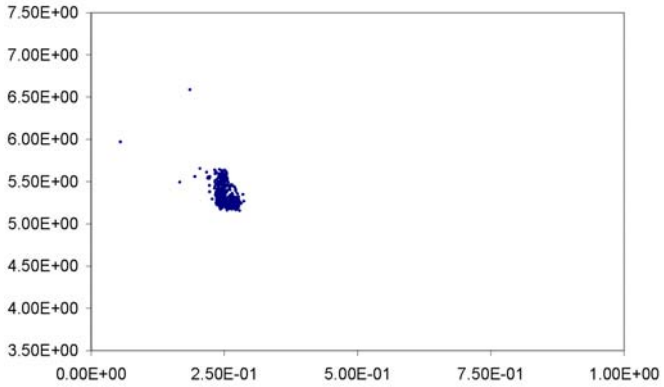


Figure 29: λΔ=0.30 $\mu_{min}/f_{\mu}$ scatter-plot; format as for Figure 24; data from the 500 repetitions illustrated in Figures 17 and 23.

$\mu_{min}$ is interpreted as a *percentage* margin. Figures 30 to 35 show these normalised data. The normalised data shows that Mode A is characterised by $\mu_{min}/\lambda\Delta>1.0$, $16<f_{\mu}/\lambda\Delta<18$; while Mode D is characterised by $\mu_{min}/\lambda\Delta<1.0$, $17<f_{\mu}/\lambda\Delta<19$. And so it is clear that the dominant mode for λΔ=0.30 is one that acts as only a very weak attractor for λΔ=0.25 (with the hindsight gained by this analysis, Mode D can now be seen as the light scattering of nine data-points above Mode A between $Q_s$=0.375 and $Q_s$=0.500 in Figure 18).

Having clarified the nature of the four modes seen in the fitness data of Figures 12 to 17, we can turn our attention to the nature of the best mode in each experiment. For $\lambda\Delta$ in the range [0.26,0.30] we see in Figures 13 to 17 that, for both Modes A and D, there is a clear shallow "U" curve with its minimum $f_\mu$ (i.e., most desired market dynamics) at $Q_s$ values around 0.125. This is true also of the Mode A curve for $\lambda\Delta$=0.25, but in that instance we see also that Mode C actually gives lower fitness (i.e., better market dynamics). Given the upward path of Mode C identified as $\lambda\Delta$ increases, we can expect to see Mode C acting as a much stronger attractor for $\lambda\Delta$<0.25. Given the large differences in genome $\mu_{min}$ and optimal $Q_s$ values between Mode C and Modes A/D, it would appear that $\lambda\Delta$=0.25 marks the point of a "phase transition" between two radically different modes of elite genotypes. To confirm the existence of this transition, two more sets of 500 experiments were conducted, for $\lambda\Delta$=0.24 and $\lambda\Delta$=0.23. Figures 36 and 37 show the respective $Q_s/f_\mu$ scatter plots, and Figures 38 and 39 show the respective $\mu_{min}/f_\mu$ scatter plots.

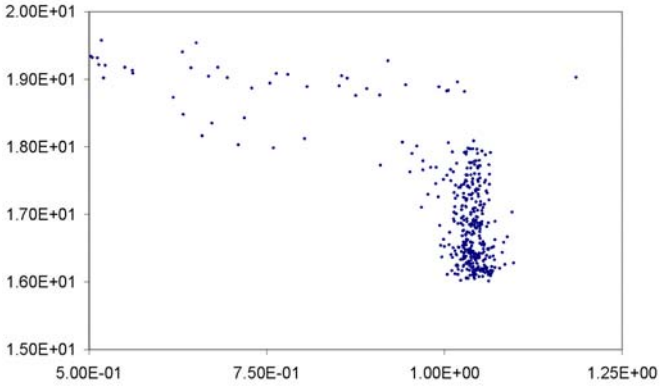


Figure 30: $\lambda\Delta$=0.25 $\Delta\lambda$-normalised $\mu_{min}/f_\mu$ scatter-plot. Horizontal axis is quotient $\mu_{min}/\Delta\lambda$; vertical axis is quotient of mean elite fitness and $\Delta\lambda$. Data from the 500 repetitions shown in Figures 12, 18, and 24.

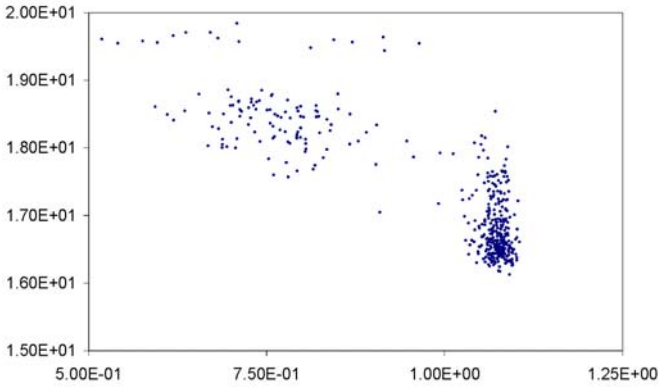


Figure 31: $\lambda\Delta$=0.26 $\Delta\lambda$-normalised $\mu_{min}/f_\mu$ scatter-plot; format as for Figure 30; data from the 500 repetitions illustrated in Figures 13, 19, and 25.

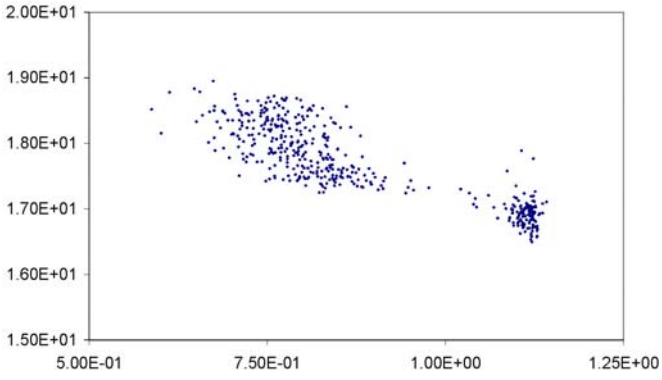


Figure 32: $\lambda\Delta$=0.27 $\Delta\lambda$-normalised $\mu_{min}/f_\mu$ scatter-plot; format as for Figure 30; data from the 500 repetitions illustrated in Figures 14, 20, and 26.

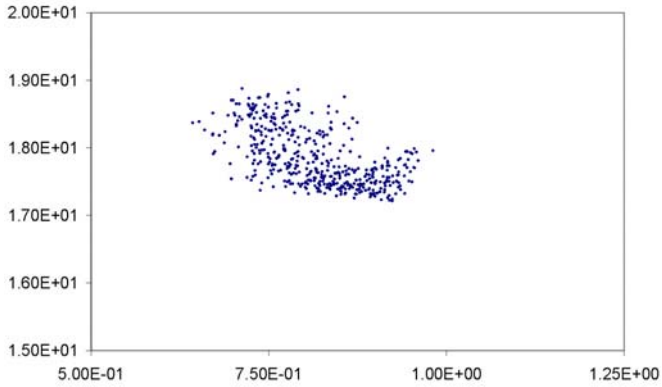


Figure 33: $\lambda\Delta$=0.28 $\Delta\lambda$-normalised $\mu_{min}/f_\mu$ scatter-plot; format as for Figure 30; data from the 500 repetitions illustrated in Figures 15, 21, and 27.

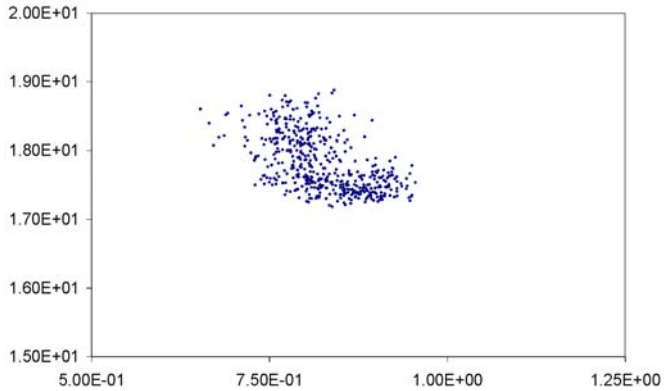


Figure 34: $\lambda\Delta$=0.29 $\Delta\lambda$-normalised $\mu_{min}/f_\mu$ scatter-plot; format as for Figure 30; data from the 500 repetitions illustrated in Figures 16, 22, and 28.

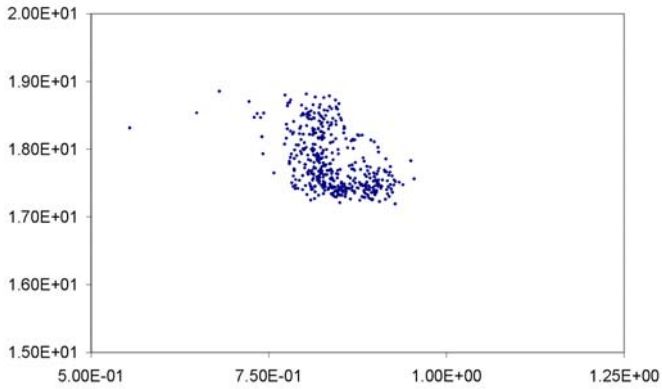


Figure 35: $\lambda\Delta$=0.30 $\Delta\lambda$-normalised $\mu_{min}/f_\mu$ scatter-plot; format as for Figure 30; data from the 500 repetitions illustrated in Figures 17, 23, and 29

The increased dominance of Mode C as λΔ reduces is clear from these extra data. The sudden transition of the elite from Mode A/D to Mode C is illustrated in Figures 40 and 41, which respectively show the $f_\mu$ values and the elite-mode $\mu_{min}$ values, each as a function of λΔ, for the best-scoring 1% ($n$=5) of the 500 repetitions for each value of λΔ. While Figure 40 shows a progressive increase (i.e., a worsening) in the fitness of the best 1% as λΔ increases, Figure 41 illustrates the "phase transition" step-change in $\mu_{min}$ values at λΔ=0.25.

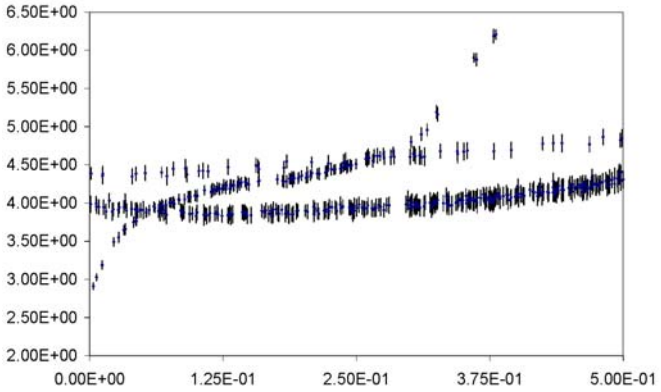


Figure 36: λΔ=0.24 $Q_s/f_\mu$ scatter-plot; $n$=500.

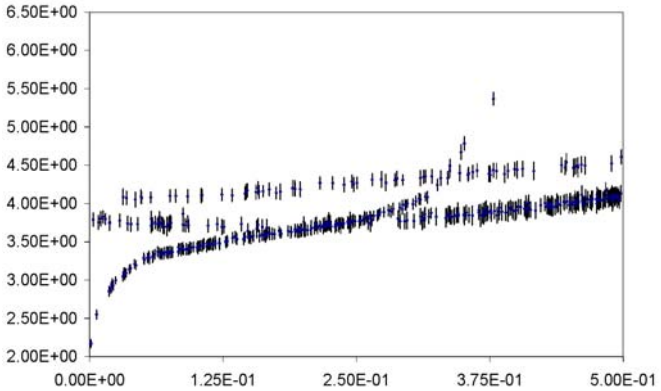


Figure 37: λΔ=0.23 $Q_s/f_\mu$ scatter-plot $n$=500.

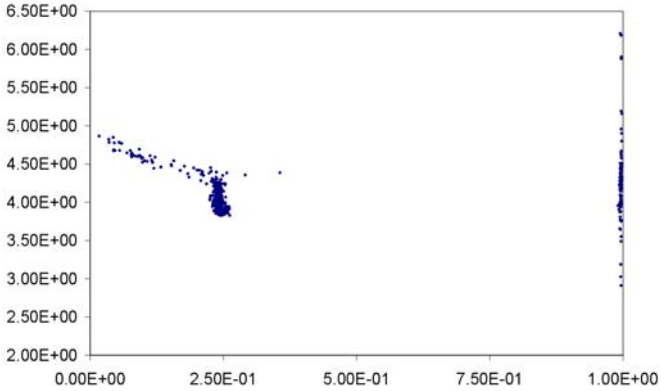


Figure 38: λΔ=0.24 $\mu_{min}/f_\mu$ scatter-plot; format as for Figure 24; data from the 500 repetitions illustrated in Figure 36.

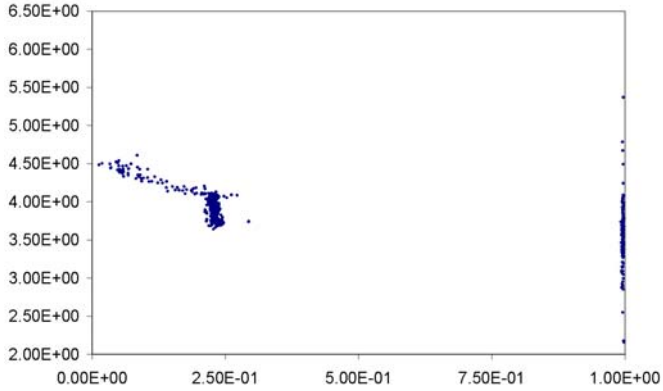


Figure 39: λΔ=0.23 $\mu_{min}/f_\mu$ scatter-plot; format as for Figure 24; data from the 500 repetitions illustrated in Figure 37.

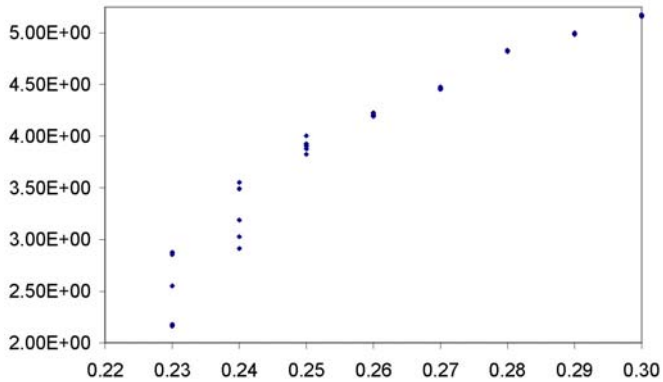


Figure 40: λΔ/$f_\mu$ scatter-plot: horizontal axis is λΔ, vertical axis is $f_\mu$; data-points are from the best-scoring 1% ($n$=5) of the 500 repetitions performed for each value of λΔ.

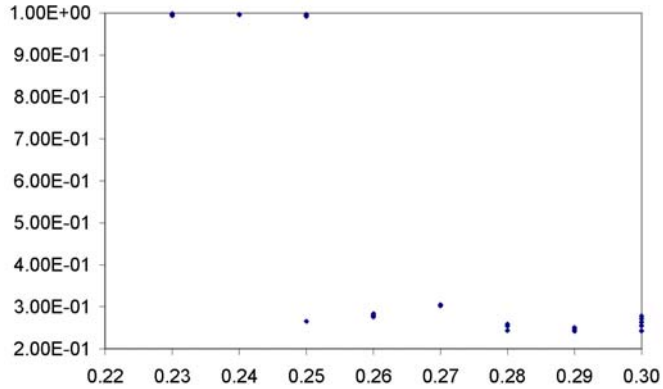


Figure 41: λΔ/$\mu_{min}$ scatter-plot: horizontal axis is λΔ, vertical axis is $\mu_{min}$; data-points are from the best-scoring 1% ($n$=5) of the 500 repetitions performed for each value of λΔ.

### *C. Discussion*

The results presented in this section, from 4000 repetitions of the GA experiments, are more illuminating than the previous contour-plot visualizations of these fitness landscapes first published in [11]. Four distinct elite-genotype modes have been identified, and the effects of changes in λΔ on those modes have been revealed. In several places, a "snapshot"

approach, i.e. studying $f_\mu$ distributions for a single fixed value of $Q_s$, (as used in the contour plots of Figures 7 to 11) would give apparently bi-modal results where in fact the data is tri-modal but two modes overlap in the $Q_s/f_\mu$ scatter-plot: this is seen most clearly in Figure 37 where Modes A and C overlap for much of the range $Q_s$=0.125 to $Q_s$=0.250.

While it is fortunate that the four modes reveal themselves so clearly as distinct clusters in the $\mu_{min}$ data, if this had not been the case then any appropriate clustering algorithm could be used to automatically identify such clusters (see [13] for a comprehensive review of data-clustering algorithms). Similarly, it is fortunate that the single $Q_s$ dimension yielded such revealing scatter-plots. But in the absence of such good fortune, there are many techniques for dimensionality reduction in multivariate data (such as principal components analysis and multidimensional scaling [14] or "principal curve" techniques such as nonlinear principal component analysis [15]) that could be employed to identify the best dimensions to use as basis vectors for projecting down from the high-dimensional source data to the two or three dimensions appropriate for display on screen or on paper.

The predominant change resulting from successive reductions in $\lambda\Delta$ from 0.30 to 0.23 is the increasing influence of Mode C as an attractor for elite genotypes. The visualization of the search space has demonstrated that $\lambda\Delta$=0.25 marks the point of a phase transition where the fittest genotypes found in an evolving-$Q_s$ GA experiment cease to come from Modes A/D and start to come from Mode C. At face value, this would indicate that Mode C genotypes are the best to use for markets such as M1 where $\lambda\Delta$<=0.25. However, this is not necessarily the case. In fact what the $\mu_{min}$ visualizations reveal is an unanticipated exploitation of a system parameter by the GA. Put more directly: Mode C genotypes score a high fitness by exploiting an aspect of the ZIP implementation not considered in the fitness evaluation function. Recall that earlier in this paper (at the end of Section II.A) it was stated that in the original ZIP implementation published in [1] (and used in all successive ZIP work: [5,6,7,8,11]) there is a system parameter MAX_FAILS, set to 100, that determines the maximum number of successive ignored quotes allowed to pass before a trading period is ended. This parameter prevents ZIP marketplaces from continuing indefinitely with traders constantly quoting prices ignored by the contraside (and making small adjustments to their margins as a consequence) but never actually trading.

The Mode C genotypes exploit the fact that, with MAX_FAILS=100, very long sequences of quotes can occur without any transactions taking place: note that the Mode C genotypes give their best performance (lowest $f_\mu$ values) at $Q_s$=~0.0 for all the values of $\lambda\Delta$ explored here. That is, in practice, the best Mode C genotypes do not allow the sellers to quote *at all*. Note also that the Mode C genotypes are distinguished by their common feature of $\mu_{min}$=~1.0. This determines that all buyers and all sellers will have initial profit margins very close to 100%. As the sellers cannot quote, the seller margins are not particularly relevant. However, with buyer margins close to 100%, at the start of trading the bid-prices in the market will all be very close to zero. Thus the market's *apparent* demand curve (i.e., the demand curve that could be constructed from the bid array given by all buyers' current $p_i(t)$ bid values, as opposed to the *true* demand curve constructed from the buyers' private $\lambda_i$ prices) will lie very close to the zero line, and it will not intersect with either the true or the apparent supply curve. All the bid prices will then gradually rise as the buyers reduce their margins each time a bid is ignored by the sellers; this gradually raises the apparent demand curve, until an intersection with the apparent supply curve does occur, at which point a transaction will take place. Significantly, the seller in each such transaction is most likely to be the current most intra-marginal seller (i.e. the currently active seller whose $\lambda_i$ value lies below the theoretical equilibrium price by the greatest amount). In this way, the evolved $Q_s$=0.0/$\mu_{min}$=1.0 mechanism ensures that extra-marginal sellers (i.e. those whose $\lambda_i$ values lie above the theoretical equilibrium price) are *highly* unlikely to engage in any transactions. A marketplace where extra-marginal sellers are effectively excluded from the outset is likely to show faster equilibration than one in which the extra-marginal sellers can trade, which helps explain why $Q_s$=0.0/$\mu_{min}$=1.0 genomes score well on the fitness measure used here.

In summary then, for Mode C the bids announced in the market during the protracted periods of non-trading allow the traders to adjust their margins such that, *when transactions do start to occur*, equilibration is significantly quicker. And so we see that the Mode C data is a reflection of the old adage "you get what you measure". The fitness evaluation function used here (and in all our previous work) is dependent only on *transaction* prices, and totally ignores the *number of ignored quotes* in the market leading up to each transaction. This is exploited by the best Mode C genotypes, which specify ZIP initialisation parameters that collectively ensure a high likelihood of *both* many failed bids preceding a transaction *and* a bid actually being accepted before the MAX_FAILS limit is reached, where the seller accepting the bid is likely to be the most intra-marginal seller remaining in the market.

One obvious way around this type of solution would be to employ a fitness evaluation function that rewards higher ratios of accepted-to-ignored quotes as well as rewarding speed of equilibration. As these are two distinct objectives, a Pareto-optimising approach would be appropriate (e.g. [16]).

Although the results presented in this paper have illuminated the nature of the surprising difference in the fitness-landscape contour-plots for M1 and M3 shown in Figures 7 and 9, and although the discussion presented in this section has identified how the solution encoded on Mode C genomes operates, several significant open questions remain to be answered.

These questions include:

- Why do Mode C genomes dominate (i.e., give the best market dynamics) for $\lambda\Delta<=0.25$ while Modes A/D dominate for $\lambda\Delta>0.25$?
- Why is there the transition from Mode A to Mode D as $\lambda\Delta$ increases from 0.26 to 0.28, and to what extent if any is that transition significant?
- Why do both Mode A and Mode D show (in Figures 18 to 23, and also Figures 36 and 37) a clear U-curve with a low-point optimum at $Q_s$=~0.125?
- What other elite-genome modes exist (if at all) for the so-far unexplored ranges $\lambda\Delta<0.23$ and $\lambda\Delta>0.30$?
- To what extent are the data presented here dependent on other free parameters in the market schedules? For instance, the range $0.23<=\lambda\Delta<=0.30$ has been explored in all the experiments reported here with 11 traders on each side, and with an equilibrium price of 2.00, and with an equilibrium quantity of 6 (see Figures 1 and 3). At this stage it is not at all obvious how the results presented here might alter quantitatively or even qualitatively as the $\lambda\Delta$ values are held constant while these other market supply/demand parameters are altered.
- What happens when the $\lambda\Delta$ on the supply curve and the $\lambda\Delta$ on the demand curve are not equal (i.e., when the slopes of the supply and demand curves are *not* symmetric)? Or when the values are not uniform for each curve (i.e., when the supply and/or demand curves are non-linear?). That is, we should explore all of the above questions for a much wider range of styles of supply/demand schedule; as Walia [11,12] has recently done for evolving-$Q_s$ markets populated by ZI-C traders.
- In exactly what way (if at all) do the results presented here depend on the values of system parameters such as MAX_FAILS or on implementation details such as the particular genetic encoding used?

All of these questions appear to be fertile ground for future research, and for the time being the most productive way of tackling these issues appears to be by empirical exploration methods such as those presented here. The results shown in this section make clear that small variations in one free parameter (i.e., the $\lambda\Delta$ value affecting the simple stationary symmetric supply and demand schedules in the ZIP-trader marketplaces) can have effects that are not intuitively predictable in advance. This lack of an intuitive understanding could in principle be addressed by appropriate mathematical modelling, but the compounded non-linearities, stochasticity, and adaptation (giving non-stationary probability distributions) in the ZIP-trader system all combine to make a full analytic understanding (e.g., via game-theoretic analysis or probabilistic modelling) seem a very distant goal. The chicken-and-egg circularity in the relationship between theory and data is well known; for the time being, we are concentrating on building a body of parametric-variation data that can then be used to suggest and constrain hypotheses/theories that guide the direction of subsequent empirical studies.

Clearly, there is no shortage of topics for further research; and nor is there any shortage of experiments with which to keep our compute-clusters busy.

## V. CONCLUSION

This paper commenced with a review of previous experimental work up to and including the recent "low-resolution" contour plots that give meaningful 2-d projections of the 10-d fitness landscapes underlying this work. The major new results of Section IV showed the outcome of 4000 repetitions of our GA experiments (taking ~8000 hours of computation).

This higher-resolution data showed that, with only very minor variations in the $\lambda\Delta$ value, the "attractor" modes for the elite genotypes in our experiments could show major shifts and phase transitions. While the $\lambda\Delta$ value is not an *explicit* term in the fitness evaluation function, it is an *implicit* factor insofar as it determines the supply and demand curves that underlie the equilibration-based fitness evaluation function used in this work; and of course the evaluation function plays a large part in determining the fitness landscape for any GA system.

The new hi-resolution data again confirms that the GA can successfully find optimum genotypes. However, the data on *sub-optimal* solutions (i.e. the nature of the fitness landscape at points *other* than optimum) has the potential to be more informative than the optimum data taken alone. Although none of the evolved solutions here improve significantly on any of our original GA results presented in [6], the new data helps us to better illuminate the nature of the fitness landscapes traversed by the GA in attempting to evolve new mechanism designs for ZIP-trader marketplaces, and also helps us to understand the nature of the optima in those landscapes. With this illumination comes the chance of better understanding *what* factors in the system affect the nature of the optimum solutions, and *how* they do so: at the moment, the causal mechanistic interactions leading to one genome being better than another for any particular supply/demand schedule remain unclear, rendering all our results as idiopathic. But the relationship between $f_\mu$, $\mu_{min}$ and $\lambda\Delta$ demonstrated in Figures 30 to 35 appears to be one promising lead in identifying such causes.

While the high-resolution visualisations presented for the first time here offer better opportunities for understanding the performance of the evolving-marketplace ZIP-trader systems, it should be noted that none of the figures in this paper are at all novel in terms of the nature of the visualisation techniques employed: scatter plots are nothing new.

But what is perhaps more novel here is the approach to consuming CPU cycles: burning 8000 hours of CPU time to generate these figures would have seemed an astonishingly profligate use of compute resources only a few years ago. And indeed, for a researcher working with a single-CPU machine, waiting eleven or so months for a set of possibly-illuminating visualization results would perhaps not be the best use of computer time.

Yet, with many hardware vendors now pushing to build "utility data centres" and/or "grid" compute facilities with thousands or tens of thousands of connected processors available offering CPU cycles as a utility-style resource [17], the shared 50-server compute farm used for this work starts to appear decidedly lightweight, despite the fact that it can be spoken of as being roughly equivalent to a single-CPU 90Ghz Pentium4 machine with 25Gb of RAM. Within the next few years, it seems perfectly plausible that research such as the explorations reported here could be conducted by buying CPU-cycles from a utility data centre (UDC) in such a way that the time taken to generate the visualization data is measured in minutes or even seconds, rather than days or weeks.[1] In such a scenario, the transition in working styles will be similar to the move from *batch-mode* to *interactive* computing that was brought about by the development of mini/microcomputers in the 1960s and 1970s. The need would then be not so much for newer and more efficient or clever ways of getting a GA to find a solution, but rather for faster and more informative ways to visualise the high-dimensional data streams pouring out of the UDC; and for more intuitive ways to steer through these data spaces to find areas of significance or interest. In fact, research interest in "blind" search techniques such as GAs and other forms of artificial evolutionary computation could possibly wane, to be replaced by a new style of working involving applications of brute-force enumerative search to spaces of possible solutions so large as to be considered practically infinite, where the search is guided interactively by skilled human operators interacting with sophisticated visualisation workstations (where "visualisation" could include presentation via non-visual sensory modes such as audio, force-feedback on controllers, etc). Right now this seems a realistic and exciting possibility, but only time will tell.

*Acknowledgements*

Thanks to Andrew Byde for valuable discussions. Computing resources came from the HP Labs Biologically-Inspired Complex Adaptive Systems (BICAS) research group: http://www.hpl.hp.com/research/bicas.

## *References*

[1] D. Cliff, "Minimal-intelligence agents for bargaining behaviours in market environments". Technical Report HPL-97-91, Hewlett-Packard Labs. (http://www.hpl.hp.com/techreports/97/HPL-97-91.html)

[2] V. Smith, "Experimental study of competitive market behavior". *Journal of Political Economy* **70**:111-137, 1962.

[3] D. Gode & S. Sunder, "Allocative efficiency of markets with zero-intelligence traders". *Journal of Political Economy* **101**:119-137, 1993.

[4] R. Das, J. Hanson, J. Kephart, & G. Tesauro ,"Agent-human interactions in the continuous double auction". *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-01)*, 2001. (http://www.research.ibm.com/infoecon/researchpapers.html)

[5] D. Cliff, "Genetic optimization of adaptive trading agents for double-auction markets". *Proceedings of Computational Intelligence in Financial Engineering (CIFEr98) 1998.* IEEE/IAFE/Informs, pp.252-258, 1998.

[6] D. Cliff, "Evolutionary optimization of parameter sets for adaptive software-agent traders in continuous double-auction markets". Technical Report HPL-2001-99, Hewlett-Packard Laboratories, 2001. (http://www.hpl.hp.com/techreports/2001/HPL-2001-99.html)

[7] D. Cliff, "Evolution of market mechanism through a continuous space of auction-types". Presented at *Computational Intelligence in Financial Engineering (CIFEr02)* session at *WCCI2002*, Hawaii, May 2002. Technical Report HPL-2001-326: Hewlett-Packard Laboratories, 2001. (http://www.hpl.hp.com/techreports/2001/HPL-2001-326.html)

[8] D. Cliff, "Evolution of market mechanism through a continuous space of auction-types II: Two-sided auction mechanisms evolve in response to market shocks". Presented at *Agents for Business Automation*; Las Vegas, June 2002. Technical Report HPL-2002-128: Hewlett-Packard Labs, 2002. (http://www.hpl.hp.com/techreports/2002/HPL-2002-128.html)

[9] S. Phelps, S. Parsons, P. McBurney, & E. Sklar, "Co-Evolution of auction mechanisms and trading strategies: towards a novel approach to microeconomic design". Presented at *ECOMAS 2002*, New York, July 2002. (http://www.csc.liv.ac.uk/research/techreports/tr2002/tr02004abs.html)

[10] Z.Qin, *Evolving Marketplace Designs by Artificial Agents.* MSc Thesis, Department of Computer Science, University of Bristol, September 2002.

[11] D. Cliff, V. Walia, & A. Byde, "Evolved Hybrid Auction Mechanisms in Non-ZIP Trader Marketplaces" Technical Report HPL-2002-247, Hewlett-Packard Laboratories, 2002. (http://www.hpl.hp.com/techreports/2002/HPL-2002-247.html)

[12] V. Walia, *Evolving Market Design.* MSc Thesis, University of Birmingham School of Computer Science, September 2002.

[13] A. K. Jain, M. N. Murty, & P. J. Flynn, "Data Clustering: A Review". *ACM Computing Surveys* 31(3):264-323, September 1999.

[14] C. Chatfield & A. J. Collins, *Introduction to Multivariate Analysis.* Chapman and Hall, 1980.

[15] M. A. Kramer. "Nonlinear principal component analysis using autoassociative neural networks". *Journal of the American Institute for Chemical Engineers*, 37:223-243, February 1991.

[16] E. Zitzler, *Evolutionary Algorithms for Multiobjective Optimization: Methods and Applications.* PhD Thesis, Swiss Federal Institute of Technology (ETH) Zurich. TIK-Schriftenreihe Nr. 30; Diss ETH No. 13398, Shaker-Verlag Germany, December 1999. (ftp://ftp.tik.ee.ethz.ch/pub/people/zitzler/Zitz1999.pdf)

[17] Information on utility clusters and grid computing can be found at www.gridforum.org, www.gridcomputing.com, www.ieeetfcc.org, and http://www.hpl.hp.com/about/vision.html#infrastructure.

[18] N. Robinson, *Evolutionary Optimisation of Market-Based Control Systems for Resource Allocation in Compute Farms*. MSc Thesis, University of Sussex School of Cognitive and Computing Sciences, September 2002.

[1] Incidently, a recent masters thesis by Robinson [18] replicates the results in [6] and explores the use of evolved-$Q_s$ ZIP-trader marketplaces for market-based control of load-balancing in minimal abstract models of UDCs.