Evolving Market Design

Vibhu Walia
Intelligent Enterprise Technologies Laboratory
HP Laboratories Bristol
HPL-2002-289
October 30th , 2002*

automated market design, ZI traders, electronic marketplaces, evolution strategies, BICAS

The Continuous Double Auction (CDA) is one of the most popular of all auction markets in the world. Some of the biggest trade markets in the world including the New York Stock Exchange (NYSE) and the Chicago Mercantile Exchange are organised as CDAs. Previous work by Cliff has shown that previously unknown 'hybrid' variants of the traditional CDA can be found to give the most desirable market dynamics. Cliff's results were based on experiments conducted using a computational simulation of the CDA populated by electronic Zero Intelligence Plus (ZIP) traders and his work uses a GA to co-evolve the market mechanism as well as the ZIP agent parameters. The exclusive use of the ZIP trading algorithm in all his work raises questions about the robustness of his results to any change in the trading algorithm used. In this thesis we use a self-adaptive Evolutionary Strategy (ES) to explore the space of possible auction types in a CDA populated by Gode and Sunder's cognitively simple Zero Intelligence (ZI) traders. We show that hybrid CDAs are still preferred over traditional variants and our results provide the first demonstration that hybrid variants of the CDA can provide favourable dynamics for trading strategies other than ZIP.

* Internal Accession Date Only

The University of Birmingham

School of Computer Science

MSc in Natural Computation

Summer-Project

Evolving Market Design

By:
Vibhu Walia

Supervisor:
Prof Xin Yao
School of Computer Science
University of Birmingham

Industrial Co-Supervisors:
Dr Andrew Byde
Prof. Dave Cliff
Biologically Inspired Complex Adaptive Systems (BICAS) research group
Hewlett-Packard Laboratories, Bristol

September 2002

Table of Contents

Acknowledgment

While pursuing my MSc in Natural Computation I have been supported by a scholarship that is partly funded by the European Commission through a grant (ASI/B7-301/97/0126-08) in its ASIA IT&C programme and partly by the School of Computer Science at the University of Birmingham. Additional support for this summer-project has come from the Biologically Inspired Complex Adaptive Systems (BICAS) research group at Hewlett-Packard Laboratories, Bristol.

I would like to thank Prof. Xin Yao for his support and guidance. I would also like to thank Prof Dave Cliff and Dr Andrew Byde for their supervision and guidance. Their help and co-operation helped make my stay at Hewlett-Packard Labs an extremely enjoyable and productive experience.

Chapter 1

Introduction

This chapter outlines the scope of our work and presents a broad overview of the rest of the thesis.

1.1 Overview

Although the Continuous Double Auction (CDA) is one of the most popular types of markets in the world today, it is also one of the least understood. Its complex nature means that game-theoretic analysis of the CDA is intractable except for the simplest of cases. Experimental methods have long been used to explore the mechanisms behind some of its most curious properties and CDA experiments with human participants have been organised since the early 60's (Smith, 1962).

Agent-based simulations offer us a powerful tool which we can use to study the properties of a CDA using empirical methods. Agent-based methods are not only convenient and economical but also offer a tightly controlled environment in which a diverse range of experiments can be performed.

Trading-agent design for the CDA is an established field and recent work by Das et al (2001) has shown that trading-agents can consistently outperform human traders in laboratory experiments. This has led some leading researchers (Kephart, 2002; Preist 1999) to predict that trading agents will play an increasingly important role in the markets of the future.

Some recent results by Cliff (2001; 2002) have shown that hybrid variants of the CDA, which are unlike previously known versions, can give the most desirable market dynamics. Work by Phelps et al (2002) uses a Genetic Programming (GP) based approach to co-evolve a trade-settling formula for the CDA along with the trading algorithm parameters. These attempts at automatic market design are particularly relevant in the context of market mechanism design for electronic trading agents. Electronic agents are willing participants in any kind of auction and there is nothing that stops us from adopting these 'peculiar' hybrid versions of the CDA for online auction markets. However, for wide adoption of these markets we need to conclusively prove that: -

1. Automatically designed markets consistently provide larger gains compared to more traditional market mechanism types.
2. These gains are achieved regardless of the trading strategy used.

Previous work has addressed the first of these concerns but has failed to address the second.

In this work we have used a Self-Adaptive Evolutionary Strategy to search through the space of possible auction types, where markets are parameterised by a real-valued parameter **Qs.** This parameter was first introduced by Cliff (2001a) and forms the basis for his work exploring automatic market design (Cliff 2001a; 2002). We have used the Zero Intelligence (ZI) traders, first introduced by Gode and Sunder (1993), to populate our simulations of the CDA. ZI traders submit random bids subject to the constraint that they cannot shout loss-making offers and specify the lower bound on the cognitive ability that is needed for a trader participating in a CDA. We also incorporate some suggestions made by Preist and van Tol (1998) and Tesauro and Das (2001) in the methodology we use to organise the CDA.

Our results demonstrate that 'hybrid' variants of the CDA dominate more traditional variants in many cases. Thus, we tackle both the issues highlighted above and demonstrate that not only do the hybrid markets, found by the ES to be optimal, offer larger gains than traditional variants, but also that these gains can be achieved in markets populated by more than one type of trading-agent. Variations in some of the implementation details of how the actual CDA is organised (compared to Cliff (2001a; 2002)) seem to also suggest that the results are not strongly tied to a particular type of CDA. Due to the simplistic and random nature of the ZI trading algorithm, we feel that our experiments provide strong support for a claim that the results achieved are robust to any changes in trading strategy.

We go on to explore the possible reasons for the superiority of the new asymmetric 'hybrid' variants over more conventional forms of the CDA. We conduct experiments to optimise Qs values over a 2-dimensional space of test-markets parameterised by the slopes of the supply and demand curves. We plot 3-dimensional landscapes of optimum Qs values for all markets in our test space and discuss possible reasons for the regularities observed.

1.2 Organisation of Report

The rest of the thesis is organised as follows.

Chapter 2 *Background Work* begins with a discussion of some basic economics concepts that are necessary for understanding the work that follows. It explains what supply and demand curves are and how they are used to describe and characterise markets. It discusses some properties of markets and defines some performance measures that can be used to study their properties. It then outlines previous work that forms the basis of our research. Smith's pioneering experiments with human traders in the 1960's that led to the establishment of experimental economics as a discipline are discussed in brief. We then describe how advances in computing power and techniques have led to the study of markets using agent-based simulations and discuss some electronic-trading agent strategies that have been designed for the CDA, including Gode and Sunder's ZIC and Cliff's ZIP traders. The seminal experiments at IBM's T.J. Watson Research Labs that demonstrated the superiority of electronic trading agents over human traders are described in brief. They serve as motivation to our discussion of Cliff's innovative experiments that attempt to automate the process of market-mechanism design. These experiments use a representation of the CDA that is parameterised by a variable **Qs** which regulates the bidding process followed in the actual CDA. A thorough exploration of this idea forms the basis of the rest of our work.

Chapter 3 *Evolutionary Design Optimisation* describes the experiments we have conducted to optimise the value of **Qs** for any market organised as a CDA. We first discuss how our experimental set-up relates to the set-up used in earlier work including Smith (1962), Gode and Sunder (1993) and Cliff (1997; 2001a; 2002). We describe in detail the choices we have made in designing our version of a computational simulation of a CDA and the rationale behind them. We then describe the Evolutionary strategy (ES) that we have designed for optimising the Qs value for any given market. We describe results obtained by using this ES on 5 test markets. We describe the nature of the fitness landscapes for each of these cases and discuss how it affects the behaviour of the ES. We also discuss the affect of changing parameter values of the ES on its performance and reliability of results obtained.

In Chapter 4 *Gaining Insight in to Empirical Results* we try to investigate what underlying characteristics of a given market could lead to a particular Qs value being optimum. We do this by constructing a large 2-dimensional space of test markets parameterised by the slopes of the supply and demand curves and determine the optimum Qs value for each of these. We plot a 3-dimensional optimum-Qs Landscape and try to study the regularities observed. We describe the method that is used to construct our test markets and the modifications we have made to our ES to gain computational leverage in these tests. We make some qualitative observations about our results and point out how this work can be extended to gain a better understanding of the CDA as an institution as well as for designing better CDAs.

Chapter 5 *Conclusion and Directions for Future Research* provides a summary of our work and points out some directions in which this work can be extended.

Chapter 2

Background Work

This chapter describes the previous work on which our research is based.
In Section 2.1, we first briefly summarise some basic economics concepts that are important for understanding the rest of the text. We then highlight the importance of the Continuous Double Auction (CDA) as an institution and discuss some early research into its characteristics.
In Section 2.2, we discuss some trading algorithms that have been designed for the CDA. In Section 2.3, we discuss the groundbreaking Agent-Human interaction experiments conducted at IBM's T.J. Watson Research Labs that established the superiority of these electronic-traders over human-traders.
Finally, in Section 2.4, we discuss Cliff's previous work in which he proposed a novel way of designing marketplaces for electronic-traders. Thorough exploration of this idea is the basis for the rest of our work

---

2.1 Background Economics

In economics a *scarce resource* is defined precisely as: "…one for which the demand at a zero price would exceed the available supply" (Begg et al 1994). The efficient allocation of scarce resources without the presence of a central controlling authority in free-markets remains, what pioneering experimental economist V.L. Smith has called, a *"scientific mystery"* (Smith, 1982).

A *market* can be defined as "…a set of arrangements by which buyers and sellers are in contact to exchange goods or services"(Begg et al 1994). What makes the markets so interesting as an institution to economists and computer scientists alike is their self-regulating nature. This means that for a given market if the supply and demand schedules remain constant, then the resource-allocation in the market is *pareto-optimal* [1]. Moreover, this pareto-optimality emerges solely out of the decentralised and distributed, low-level,

[1] Pareto-optimal refers to a situation in which you can't make anybody better off without making someone else worse off.

self-satiating behaviour of the agents active in the market. This is a typical example of what computer scientists refer to as an *emergent phenomenon*. This seemingly mysterious phenomenon has commonly been referred to as the influence of the "*Invisible Hand*" by economists, ever since Adam Smith coined the term in the late eighteenth century (Smith, 1776).

2.1.1 Supply and Demand

A graphical representation of the supply and demand schedules in a typical market is shown in Figure 2.1 below. Each seller has a cost price c_i for each item i that she possesses, and in most cases, this is the lowest price at which she will sell item i. Similarly each buyer has a utility price u_i for each item i that she wishes to buy and this is the maximum price that she will be willing to pay for it.

The graph shows two curves each of which represent

1 The quantity that will be available for purchase by a buyer at a given price (Supply curve).

2 The quantity that is wanted for purchase at a given price (Demand curve).

As shown, the quantity available for purchase increases with price. This is because more suppliers are willing to sell for a higher price as opposed to a lower price. The opposite is true for the demand curve – more buyers are willing to purchase a good at a lower price.

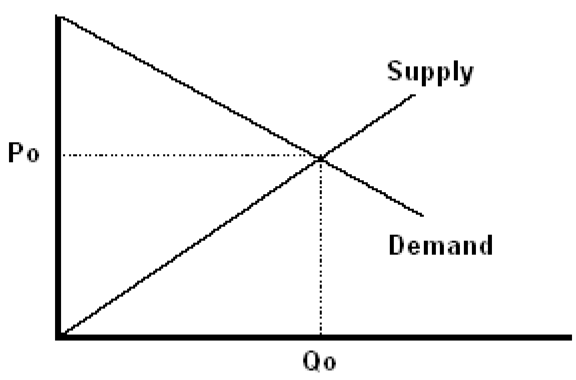

Fig 2.1. A simple Supply and Demand schedule. The intersection of the supply and demand curves gives the Equilibrium Price (P_o) and Equilibrium Quantity (Q_o). See text for discussion.

Equilibrium Formation

Given a stable underlying supply and demand, the transaction price in the market will converge to a clearing or *Equilibrium Price* ***Po***, at which the quantity supplied is equal to the quantity demanded. The value of the equilibrium price and quantity can be determined from the graph of supply and demand schedules – the point of intersection of the supply and demand curves gives us the equilibrium price and quantity. This is shown graphically in Figure 2.1.

But more precisely, the *Equilibrium Quantity* ***Qo*** is the maximal number *n*, for which there exist *n* sellers, *n* buyers, and a price *Po* such that all of the buyers and sellers are willing to trade at *Po*. Any such price *Po* is the equilibrium price.

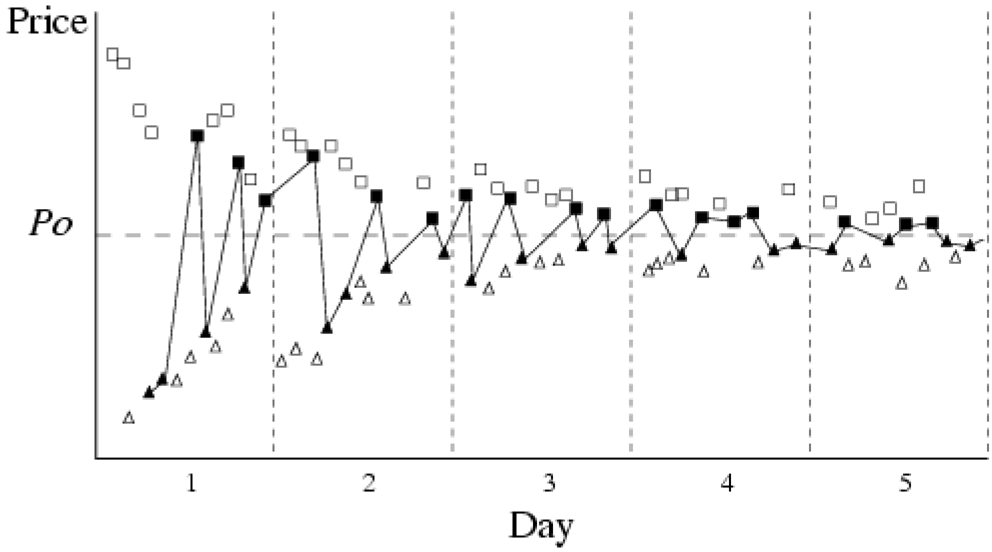

Fig 2.2. Time series for both accepted and rejected bids and offers (taken from Cliff (1997, Figure 2, p. 17)). This synthetic data shows how the bid-offer spread reduces over a 5-day experiment. Squares represent bids shouted by buyers and triangles represent asks shouted by sellers. The hollow symbols represent rejected quotes and the solid symbols represent the quotes were accepted. The accepted bids and offers are joined by a line which shows that the transaction price approaches the equilibrium as the experiment proceeds.

It is important to realise that in most cases, the true nature of the supply and demand is schedule is known neither to the traders nor central authority (if any). Traders simply try to maximise their profits; which for buyers is the difference between the transaction price and their true utility for the purchased good $(p_i - u_i)$ and conversely, the difference between transaction and cost price for the sellers $(p_i - c_i)$. Buyers and sellers can be imagined as pulling the average transaction price in opposite directions – equilibrium formation is a phenomenon that results from the interaction between these two opposing forces

2.1.2 Measuring Performance

The maximum possible profit, aggregated across all traders in the market, which can be made in a given market, is known as the *total market surplus*. All traders who can participate in a trading session in which the maximum available surplus in a given market is extracted are known as *intra-marginal* traders. Those who cannot are known as *extra-marginal* traders. The concept of intra and extra-marginal traders and their influence on efficiency is illustrated in Figure 2.3.

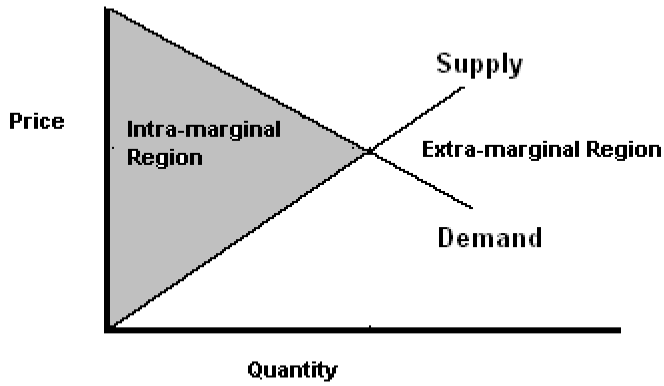

Fig 2.3. If all trades in a market are between intra-marginal traders: the market efficiency is 100%. Efficiency is reduced when if an extra-marginal trader is able to execute a transaction.

a) Allocative Efficiency

The ratio of the total market surplus to the surplus extracted in a given trading session, expressed as a percentage, gives the *allocative efficiency* for that trading session. In other words, allocative efficiency is the total profit earned by all the traders in the market divided by the maximum possible total profits that could have been earned by all the traders and expressed as a percentage.

We have exclusively used allocative efficiency as our fitness measure in all our experiments. This has to do with our choice of using ZI-C agents. ZI-C agents do not always equilibrate but can successfully extract a high-percentage of the surplus available in the market. This issue is discussed further in Section 2.2.1 when we discuss the ZI-C trading algorithm. In the rest of the thesis when we refer to efficiency, we are referring to allocative efficiency unless mentioned otherwise.

b) Smith's Coefficient of Convergence α

The equilibrating property of markets has been of particular interest to earlier researchers including Smith (1962, p. 116) who notes: -

> "The most striking general characteristic of [these] tests…is the remarkably strong tendency for exchange prices to approach the predicted equilibrium price for each of these markets. As the exchange process is repeated…the variation in exchange prices tends to decline, and to cluster more closely around the equilibrium."

To measure the convergence of the transaction prices to Po Smith defined a 'coefficient of convergence' α which is measured for each trading period. Smith (1962, p. 116) defined α as follows: -

> "The α for each trading period is the ratio of the standard deviation of the exchange prices, σ_o, to the predicted equilibrium price, Po, the ratio being expressed as a percentage. That is $\alpha = 100\ \sigma_o$ / Po"

Smith's α has been used to measure trading-agent performance in most of the previous work in trading agent-design including Gode and Sunder (1993), Cliff (1997) and Gjerstad and Dickhaut (1998). Cliff also uses α to measure the fitness of different market types in his work on automatic market design (Cliff 2001a; 2002).

However, although α measures how well agents approach the predicted equilibrium price, we feel other measures like allocative efficiency and profit dispersion can reflect our true motives behind designing better agents and markets more directly. Das et al (2001) [2] and Phelps et al (2002) exclusively use allocative efficiency in their experiments.

[2] Since Das et al (2001) study trader performance, they have use a modified definition of allocative efficiency which measures how well a trader is able extract the available surplus in the market. This modification can be seen as localising the scope of efficiency measured, from the original market-wide definition, to a trader-specific measure.

c) Profit Dispersion

Profit dispersion is the root mean squared definition between the actual profits and the equilibrium profits of individual traders. Gode and Sunder (1993, p. 133) define Profit dispersion as follows: -

> "Let a_i and π_i be the actual and theoretical equilibrium profits of traders i $\forall$ i=1, …, n. Then profit dispersion is given by $[n^{-1} \sum_i (a_i - \pi_i)^2]^{0.5}$."

Except Gode and Sunder (1993) we do not know of any other work that uses profit dispersion to as a measure either trading-agent or market-design performance. However profit dispersion can be used as a strong indicator of the social-welfare aspect of market-design.

2.1.3 Market Based Control

Wellman (1995) argues that this self-regulating nature of markets in which pareto-optimal resource allocation can be achieved in a distributed fashion makes the institution of a market ideally suited for application to most artificial intelligence problems.

He puts forth the view that most, if not all, AI problems can be reduced to equivalent resource allocation problems, which can then be solved efficiently, in a decentralised fashion by following a market-based approach. Clearwater (1995) has compiled a collection of papers that discuss the application of *Market Based Control (MBC)* to diverse areas such as bandwidth and memory allocation, job-shop scheduling and energy-conservation. Cliff (1997; 1998) provide a survey, as well as a critique, of some current applications of *MBC* to practical problems. They suggest improvements that will help truly unleash the potential of market-based systems as decentralised and automatic problem solving technique.

2.1.4 Market Organisation

Trades in actual markets (in the field or the laboratory) is governed by a *market institution,* that is, a set of rules specifying which sorts of bids and other messages are legitimate and how and when specific traders transact, given their chosen messages. The origins of these market-institutions are not well documented but can be traced back to "haggling" in marketplaces of ancient Egypt and Mesopotamia. (Arthur, 1988). Some well-known and prevalent types of markets are the *Dutch Auction, English Auction, Vickery Auction and the Continuous Double Auction.*

In the *descending offer or Dutch-Flower Auction (DFA)*, the buyers sit silent as sellers shout increasingly decreasing bids. The buyers can step in at any time and accept the latest bid[3]. In the *English or ascending bid Auction* the opposite is true. The buyers repeatedly shout increasing bids till all possible transactions have occurred.[4]

Economists have long appreciated the impact of the auctioning procedure followed on the final outcome. Game theoretic analysis has commonly been used to derive the expected outcomes for a given type of auctioning mechanism. Varian (1995) points out that the English auction will award the item to the buyer with the highest utility value, but at a price that is equal to the reserve price of the buyer with the second highest utility value (plus a small amount to break the tie).

In a *Vickrey Auction*, each participant submits a single sealed bid and the highest (in case of buyers) bid is awarded the right to transact, but at the value of the second-highest bid. It can be shown by game-theoretic analysis that the optimal strategy in such an auction is for each trader is to bid his or her true value for the goods being auctioned.

In a *Double Auction (DA)*, both buyers and sellers can submit bids and offers. The DA is a general name for a broad class of trading institution. Friedman (1993) provides a good survey of the many variants of the DA that exist in the world today.

Continuous Double Auction (CDA)

By far the most popular form of the DA is the *Continuous Double Auction (CDA)* or *open-outcry* market. In a CDA buyers and sellers simultaneously and asynchronously announce bids and offers: and at any given time a buyer is free to accept to an offer made by any seller and vice-versa. A closely related mechanism is the *clearinghouse* or *call market* in

---

[3] This kind of auction is common in wholesale markets. It gets its name from the bidding process that is followed in wholesale flower markets in Holland.

[4] The word 'auction' is most commonly associated with the English Auction as this is the procedure followed by most government institutions and art houses.

which traders repeatedly submit bids and asks to a clearinghouse, which periodically crosses bids and asks to determine a market clearing price.

The CDA has been widely studied and is used in many major institutions including commodity-trading pits like those at the Chicago Board of Trade, and its variants are used in almost all major stock exchanges around the world including those at London, Tokyo, Frankfurt as well as the New York Stock Exchange (NYSE). Even though the CDA is by far the most common type of market mechanism used around the world today, we do not understand several of its basic characteristics. Its rapid convergence to equilibrium in the absence of any external control is its most desirable characteristic but also one of its least understood properties. Metaphors like "invisible hand" and "scientific mystery" have been used to refer to this phenomenon but how a group of selfish, non-cooperating traders reliably drive towards pareto-optimal allocation in a distributed and asynchronous framework is still a source of amazement.

Standard economic theory that tries to explain the rapid equilibrium formation observed in the CDA is built upon two assumptions (Gode and Sunder, 1993): -

a) The institution of a *Walrasian tatonnement*[5], which is an auction that is conducted by a centralised auctioneer, and
b) The utility-maximising (perfectly rational) behaviour of the traders in the market

We now know that neither of these is necessary to produce high-market efficiency and this knowledge gives a strong impetus to research that is aimed at discovering the true reason for high efficiency of markets as an institution.

Using Agents to Study the CDA

A game theoretic analysis of the CDA is intractable except for the most trivial cases in which strong assumptions have to be made about both trader behaviour and the nature of the market itself. Hence, there are few other options than to either perform empirical experiments with human beings or simulate trader behaviour with the help of agents. Besides the obvious expense and effort that is required to conduct any large-scale human study, there are several factors that make agent-based empirical methods suitable for the

[5] Tatonnement is French for groping.

study of markets in particular. LeBaron (1998) discusses some of these issues, chief of which are: -

a) Issues of price and information aggregation tend to be sharper in financial settings where agent objectives tend to be clearer.
b) Accurate financial data is readily available at many frequencies ranging from a minutes by minute to a yearly basis.
c) Continuing developments in economics and finance can be compared with agent-based experiments.
d) There are many observed empirical 'puzzles'[6] that cannot be explained satisfactorily by current-day analytical models.

We provide a literature review on the field of using agent-based simulations to gain insight into complex economic and social systems in Appendix A.

Market Design

The design and performance of technical trading rules is indeed, one of these puzzles. Market-mechanism design has so far been a highly subjective skill, more of an art than a science that requires considerable skill and experience along with a high-degree of domain knowledge. Game theory helps in gaining some analytical insight into few of the properties of markets, but analysis of the CDA is intractable except for the most trivial cases. Thus, when designing mechanisms for use in market like the CDA, predictions about the advantages of choosing one kind of market-mechanism over another are mostly limited to guesswork or speculation at best. Agent-based methods allow us to answer questions about the behaviour of proposed designs thus paving the way for automatic market-mechanism design. These models allow us to answer questions about how one market performs relative to another with respect to the various performance metrics that are used to measure market fitness including allocative efficiency, profit-dispersion and equilibrium formation. We discuss more about this new vista in Section 2.4.

---

[6] LeBaron (1998) lists some of these puzzles: premium equity puzzle (Hansen and Singleton, 1983), volatility persistence (Bollerslev et al, 1990) and performance of trading rules(Sweeney, 1986)

2.1.5 Experimental Economics

Even though the use of laboratory methods in economics to test hypothesis had been of interest since the 1930's (Cliff, 1997, p. 15) perhaps the biggest impetus to the field of experimental economics came from a series of experiments conducted by Smith (1962) over a five-year period from 1955 to 1961. We now briefly describe Smith's methodology and point out some of the contributions his seminal work made.

Smith experiments were organised as follows (Smith, 1962, p. 112). A group of human subjects were divided at random into two subgroups, a group of buyers and a group of sellers. All buyers b_i received cards containing the utility value u_i (known only to the buyer) and was asked to maximise his profit (u_i-p_i) by buying a good at price p_i in the experimental auction to be conducted. Similarly sellers received cards with cost prices c_i and were asked to maximise their profits (p_i-c_i). The distribution of limit prices determined the supply and demand curve for the market and these were used to determine the predicted equilibrium price and quantity, Po and Qo.

Traders were asked to make the following assumptions about the nature of the utility and cost price values: -

a) For buyers, buying goods at utility value is better than not buying them at all (and similarly for sellers – selling at cost price is strictly preferred to not selling at all).
b) The good being traded are commodities[7], i.e. the only difference between the different goods for sale in the market is in their price; there is no other way of distinguishing one item from another.

In the CDA conducted, all traders were allowed to shout bids or offers at any time at a price that did not violate their minimum (or maximum) reservation price. All buyers and sellers were free to accept any bid (or offer) in which case a binding contract was closed and the deal making buyer and seller dropped out of the market. As soon as the contract was made the transaction price was recorded. Smith (1962, p. 113) notes that all purchases were for final consumption and resale was not allowed. This procedure was

7 Gold, steel, wheat etc. are examples of goods that are commodities.

continued until bids and offers were no longer leading to contracts. One or two final bids or offers were made after this point and the market was officially closed after that.

Experiments were conducted over several trading days (ranging from 2 to 6 days) and the convergence of the transaction prices the predicted equilibrium price was studied using the coefficient of convergence α, which was discussed in Section 2.1.2 on Page 13.

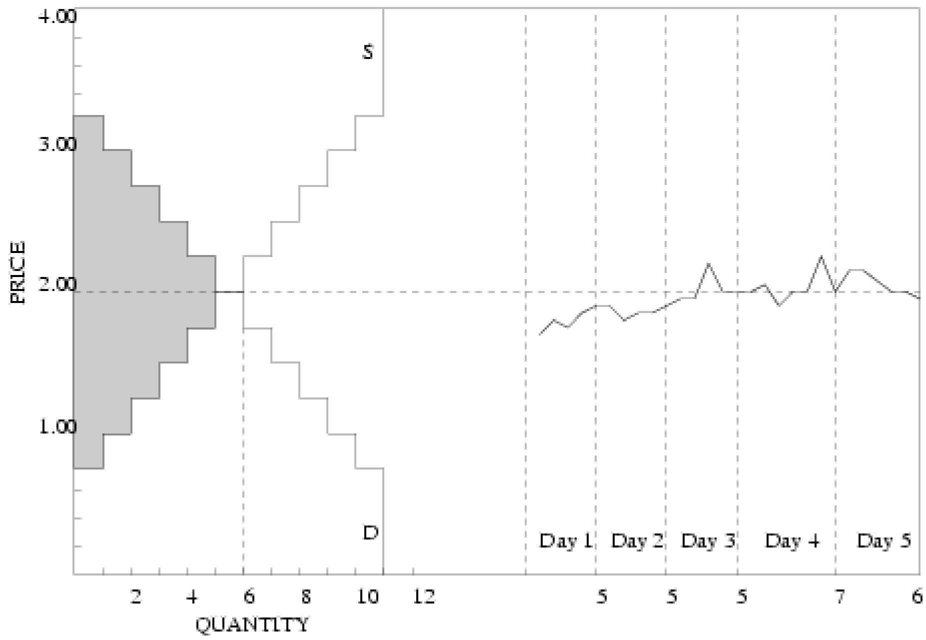

Fig 2.4. This figure shows the convergence of the transaction prices to the predicted equilibrium price Po in Smith's experiments (taken from Cliff (1997, Figure6, p.21). Orig. from Smith (1962, Chart1, p. 113)). The supply and demand curve for the market is shown on the left. The dotted horizontal line indicates the equilibrium price Po and the shaded region indicates the intra-marginal region of the market. The convergence of the transaction prices to Po is shown on the right as the experiment is conducted over a five-day period.

Results from one of Smith's original experiments are shown in Figure 2.4. The figure shows how the transaction price converges to the equilibrium price predicted from the underlying supply and demand curve. Note that all traders always shout bids that are strictly above (or equal to) their actual reserve prices. As a consequence the apparent supply and demand curve is very different from the true, underlying nature of these curves. Also, as trading proceeds, buyers and sellers leave and enter the market, causing the supply and demand schedule to change dynamically with time. So it was surprising to observe that as trading proceeded, the average transaction price rapidly approaches Po for most cases.

Besides an experimental methodology that has been subsequently adapted and used for agent-based simulation of markets by researchers like Gode and Sunder (1993), Cliff (1997), Das et al (2001) and indeed in our own work, Smith made a significant contribution by showing that neither the Walrasian tattonement institution nor a large number of traders was necessary for markets to display self-regulated equilibrium approaching behaviour. In a later publication, discussing his results Smith (1992, p. 157) states: -

> "What have I shown? I have shown that with remarkably little learning, strict privacy and a modest number [of subjects], inexperienced traders converge rapidly to a competitive equilibrium under the Double Auction mechanism. The market works under much weaker conditions than had traditionally been thought to be necessary…You do not need price-making behaviour – everyone in a double oral auction is as much a price-maker as a price-taker."

2.2 Trading Agents for the CDA

Trading-agent design for the CDA is a well-established field of research today. There has been a series of events that have provided impetus to the field starting with the seminal work in experimental-economics done by V.L. Smith, which we discussed in the Section 2.1.5. Smith's work attacked the assumed requirements for the CDA to be a Walrasian-tatonnement and to consist of an infinite (or large) number of traders and proved both of these assumptions to be false. But since Smith's experiments involved human traders it did not address the perfect-rationality assumption.

Gode and Sunder's seminal experiments (Gode and Sunder 1993), which are discussed in Section 2.2.1, were the first that separated trader-rationality from market-structure and studied both separately. These experiments established for the first time, the distinct contributions both these factors make to the overall market-efficiency. However, many of Gode and Sunder's claims were seriously flawed and this was pointed out by Cliff (Cliff, 1997) with the help of their Zero Intelligence Plus (ZIP) trading-agents, which are discussed in Section 2.2.2.

Although ZIP agents were originally developed solely to counter claims made by Gode and Sunder about the efficacy their Zero Intelligence Constrained (ZIC) traders, they received high acclaim when they consistently outperformed human traders in Agent-

Human interaction experiments performed at IBM's T. J. Watson Labs. The fact that computationally simple agents like ZIPs can consistently outperform human beings was another step forward in addressing the perfect-rationality assumption still forms the basis of many economic models.

2.2.1 Zero Intelligence Constrained (ZIC) Traders

Gode and Sunder's chief claim as outlined in the abstract of their paper (Gode and Sunder, 1993, p.119) was: -

> "Allocative efficiency of a double auction derives largely from its structure, independent of trader's motivation, intelligence or learning. Adam Smith's invisible hand may be more powerful than some may have thought; it can generate aggregate rationality not only from individual rationality but also from individual irrationality."

Indeed, this claim was well-supported by evidence given in their paper (and reproduced in our own experiments) and it was a major contribution that led to the separation of market-structure and trader-rationality as factors that lead to overall market efficiency.

Some previous work by Becker (1962) and others had tried to address the issue of individual rationality. Becker proved that several basic features of economic behaviour such as downward-sloping demand functions and upward-sloping supply functions could be derived as market-level consequences of agents' random choice behaviour subject to *budget constraints*. Here a budget-constraint refers to a rule followed by most auction markets which mandates that all traders 'settle' all outstanding trades at the end of each trading day. This means that traders must be solvent at the end of the trading day and forces them to strictly limit their behaviour to non-loss making strategies. But it does not place any other constraints on the trading-strategy they choose.

This idea seems to be the inspiration behind Gode and Sunder's work: they designed agents that exhibit completely random behaviour subject to only to a budget constraint. They cite their motivations being the reconciliation of Becker and Smith's work and thus proving that neither perfect-rationality nor a Walrasian-tatonnement assumption, are necessary for high market efficiency. In their own words (Gode and Sunder, 1993, p. 120): -

"Becker assumed that supply and demand functions yield equilibrium results through traditional tatonnement mechanisms; Smith's subjects were motivated to seek trading profits. We show that a double auction, a non-Walrasian market mechanism, can sustain high-levels of allocative efficiency even if agents do not seek profits."

Both Becker and Smith had concentrated on the equilibrium formation aspects of the markets they studied in their work. Gode and Sunder seem to have been influenced by this and even though Gode and Sunder used allocative efficiency as their primary market-performance measurement characteristic, the biggest caveat in their work arose from a desire to prove that Zero Intelligence traders could also achieve equilibrium in the double-auction market. The falseness of this claim was first pointed out by Cliff (1997).

Gode and Sunder experimented with 3 kinds of traders. The first were *Zero Intelligence Unconstrained or ZI-U traders*. These submit "…random bids and offers distributed independently, identically and uniformly over the entire feasible range of prices". The second were *Zero Intelligence Constrained or ZI-C traders*, which were the same as ZI-Us except that ZI-Cs were subject to a budget-constraint and could not shout loss-making bids or offers. Finally, they used human subjects to benchmark the results obtained with ZI-U and ZI-C agents.

Gode and Sunder emphasize that the difference between the outcomes of experiments performed with ZI-Us and ZI-Cs was attributable solely to the budget-constraint, which was an attribute not of the trading strategy, but of the market-structure. In their own words (Gode and Sunder, 1993, p. 123): -

"Traders have no intelligence in either the ZI-U or ZI-C market; the ZI-C market prevents the traders from engaging in transactions that they cannot settle. Consequently, we can attribute the differences in market outcome to the discipline imposed by the double auction on traders"

Experiments were conducted in a manner similar to Smith's original work. 12 traders were divided into 2 equal groups of buyer and sellers. Each trader is endowed with the right to sell or buy *i* units with utility (u_i) or cost values (c_i) known only to the trader. Bids, offers and transactions are valid for a single unit only and when a bid and ask cross the transaction price is equal to the earlier of the two. Each market experiment is run for six periods (or days) of specified duration and rights to buy and sell are replenished at the end of each trading period.

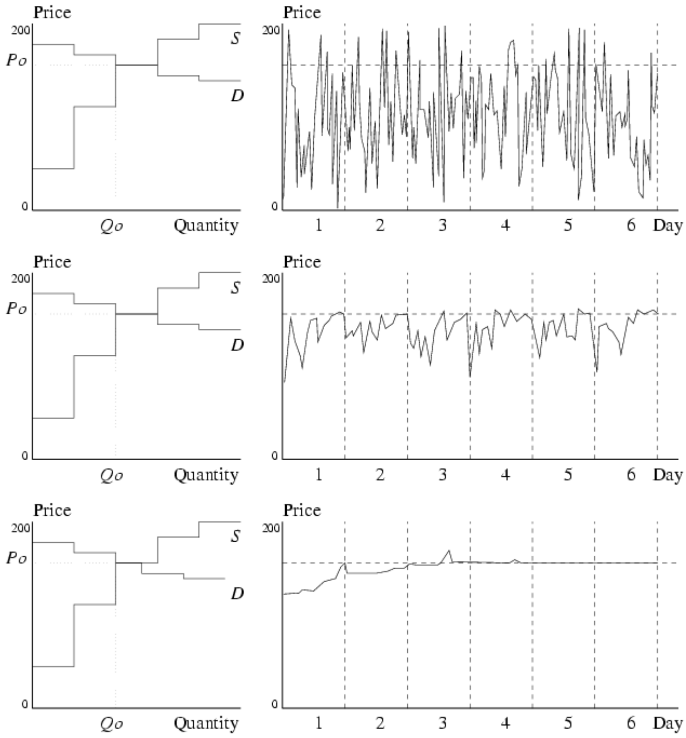

Fig 2.5 Results from one of Gode and Sunder's Experiments (taken from Cliff (1997, Fig 10, P. 25). Originally given in Gode and Sunder (1993, Fig 4, P. 127)). The Supply and Demand schedule is shown on the left for each experiment. The figures on the right show the time-series of transaction prices across a 6-day experiment. The top panel shows results for ZI-U trades, the middle panel for ZI-C traders and the bottom panel for human traders.

Gode and Sunder discuss 3 performance metrics for each of their experiments – allocative efficiency, root means squared deviation of transaction prices from equilibrium price Po (Smith's σ_o) and root mean square difference between actual and equilibrium profits (or profit-dispersion). In spite of recognising that the high allocative efficiency was being achieved by ZIC traders, Gode and Sunder repeatedly emphasize the equilibrium formation behaviour of ZI-C traders throughout their paper; this can be interpreted as an attempt to reconcile their results with the established school of thought that believes in a direct causal relationship equilibrium-formation and high market-efficiency.

They noted three features of the ZI-C price series (shown in Figure 2.5).

1 The series shows no signs of learning from period to period; the series from each period are statistically independent.
2 The volatility of the ZI-C price series is greater than those in human markets but less than ZI-U markets.
3 ZI-C price-series, though more volatile than human market price series, converges slowly towards equilibrium within each period.

Gode and Sunder (1993, p. 129) make a very strong claim about Feature 3 in: -

> "By the end of a period, the price series in ZI-C trader markets converges to the equilibrium almost as precisely as the price series from human trader markets does"

They explain this behaviour as follows (Gode and Sunder, 1993, p. 129): -

> "…this convergence cannot be attributed to learning from market participation…[it] is caused solely by the progressing narrowing of the opportunity of ZI-C traders"

Gode and Sunder put forward the hypothesis that since expected value for bids are highest for traders at the left end of the market, i.e. those with highest margins at equilibrium price, these units are traded earlier than units possessed by traders further down the market demand function. This leads to a progressive narrowing of the bid-ask spread and leads to steady convergence to equilibrium price.

Gode and Sunder present the data for root means squared deviation of transaction prices from equilibrium price Po (Smith's σ_o) for all their 5 test-cases for ZI-U and ZI-C markets and show that σ_o decreases sharply for ZI-C but not ZI-U traders (they do not provide σ_o results for human markets). They then present the allocative efficiency measures for ZI-C, ZI-U and human trader markets and successfully show that ZI-Cs are capable of achieving very high values for allocative efficiency and their performance is very similar to human traders for this measure. Finally, results for profit dispersion are given and Gode and Sunder (1993, p. 134) note that: -

> "…in contrast to aggregate efficiency, distributional aspects of market performance may be sensitive to human motivation and learning"

Gode and Sunder made a significant contribution with their work and it was a first in many respects. It successfully proved that perfect individual-rationality was not necessary for producing high market-efficiency. It also provided one of the first demonstrations of using a simple agent-based simulation to gain significant insight into a complex system – an insight that could not be gained by analytical analysis.

Perhaps their biggest failure was that they failed to realise that allocative efficiency and equilibrium formation behaviour (shown by a rapidly decreasing σ_o) are not strongly correlated and high allocative-efficiency can be achieved even in the absence of equilibrium seeking behaviour. This is discussed further in the next section.

In spite of this, Gode and Sunder made a seminal contribution with their work and their work has been, and continues to be, highly influential amongst economists and computer scientists alike.

2.2.2 Zero Intelligence Plus (ZIP)

ZIP agents were originally developed by Cliff (1997) at Hewlett-Packard Labs, UK solely to counter Gode and Sunder's claim that ZI-C agents could achieve and stabilize at equilibrium in double auction markets.

The motivation was to prove that the results concerning the equilibrium seeking behaviour for ZI-C agents was artefactual and there were markets in which ZI-Cs would fail to achieve equilibrium. There was also an attempt to redefine the lower bound for trader-rationality by designing very simple ZIP agents that could achieve and maintain equilibrium in all double auction settings. However, later research (Das et al, 2001) showed that ZIP agents could rival or exceed human-trader capabilities and thus ZIP agents provided a new lower bound on intelligence that is needed to achieve success in double auction markets.

Cliff (1997) agrees with Gode and Sunder (1993) on the claim that high levels of market-efficiency can be achieved with ZI-C agents but disagrees with the claims made about the equilibrium seeking capabilities of ZI-C traders. He states clearly that: -

"Gode and Sunder's central argument, that the structure of a double auction market is largely responsible for achieving high-levels of allocative efficiency, regardless of the intelligence, motivation, or learning of the agents in the market, is not in doubt. …the equilibrating tendencies of the ZI-C agents is questioned"

Cliff (1997, p. 27) provides using both analytical arguments and simulation results that: -

"The mean or expected value of the transaction-price distribution is shown qualitatively close to the equilibrium price only where the magnitude of the gradient of linear supply and demand curves if roughly equal"

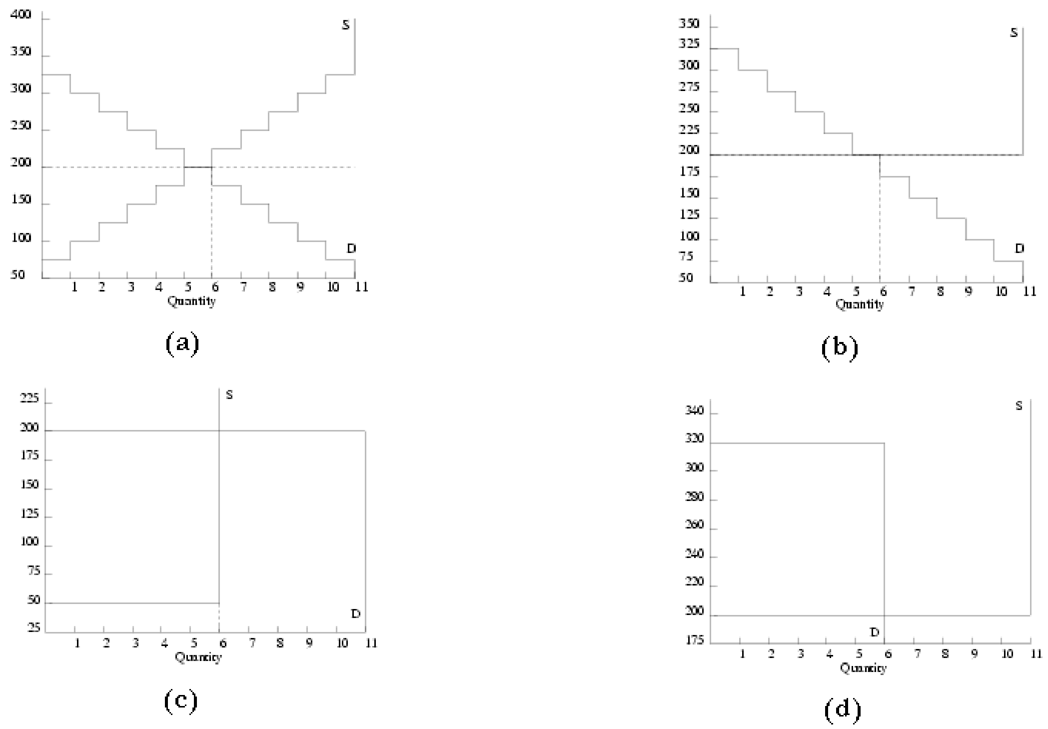

Fig. 2.6 Test cases used in Cliff (1997) to prove failure of ZI-Cs to achieve equilibrium (Reproduced from Cliff (1997, Figure 24,26,28 and 30, Pages 34-35)). ZI-Cs successfully equilibrate in market (a) but fail to do so in markets (b) to (d)

This work provided irrefutable evidence that ZI-C traders do not (and as the analytical results showed, should not be expected to) exhibit equilibrium seeking behaviour. However, at the same time ZI-C traders are capable of achieving high levels of market efficiency. This result has important implications, which one must keep in mind when choosing a performance metric for market fitness. We must be aware of the fact that the various performance metrics that have been used traditionally to measure trader and market efficacy are not necessarily correlated and a high performance on one metric does not guarantee similar performance on another. Expectations to the contrary are probably a remnant of older research, like Smith's (1962) work, which seemed to imply that metrics like deviation from equilibrium price (Smith's α and σ) and measures like allocative efficiency were tightly coupled.

Algorithm

We give a brief description of the ZIP algorithm now. For more details refer to Cliff (1997). In the ZIP strategy an agent maintains a desired profit margin $\mu(t)$ that is initialised to a random positive surplus (as in ZIC) but adjusted in a Widrow-Hoff fashion after each successful or failed trade on the basis of some simple heuristics. These heuristics are outlined in Figure 2.7. The price p(t) that a ZIP seller will shout at any time t is given by

$$p(t) = \lambda_s (1 + \mu (t))$$

where λ_s is the cost price for the seller. Conversely a ZIP buyer shouts

$$p(t) = \lambda_b (1 - \mu (t))$$

where λ_b is the utility price for the buyer

To initialise an entire ZIP trader market it is necessary to specify six market-initialisation parameters in all. Some work on GA optimisation of these parameter values is described in Cliff (1998) and Cliff (2001).

- **For SELLERS**

```
If (the last shout was accepted at price q)
    Then
        If p<=q raise profit margin
        If (last shout was a bid)
            Then
            Any active seller for which p>=q should lower its margin
    Else
        If (the last shout was an offer)
            Then
            Any active seller for which p>=q should lower its margin
```

- **For BUYERS**

```
If (the last shout was accepted at price q)
    Then
        If p>=q raise profit margin
        If (last shout was an offer)
            Then
            Any active buyer for which p<=q should lower its margin
    Else
        If (the last shout was a bid)
            Then
            Any active buyer for which p<=q should lower its margin
```

Fig. 2.7. This figure outlines the heuristics that ZIP traders use to adapt their profit margins. (Adapted from Cliff (1997, p. 43))

2.2.3 Gjerstad-Dickhaut (GD) Traders

GD agents were designed by Gjerstad and Dickhaut (1998) and recent research (Das et al 2001) has shown that they can consistently outperform human traders in laboratory experiments. We consider the GD algorithm very briefly to serve as an introduction to the Agent-Human Interaction experiments conducted by Das et al (2001) at IBM's T. J. Watson research centre that are described in Section 2.4. For more details refer to Gjerstad and Dickhaut (1998).

GD agents use the history H_M of recent market activity (the bids and asks leading to the last M successful trades) and calculate a "belief" function $f(p)$ estimating the probability or ask at price p to be accepted. For example, for a seller,

$$f(p) = \frac{AAG(p) + BG(p)}{AAG(p) + BG(p) + UAL(p)}$$

Where: -

AAG (p): is the number of bids accepted in H_M with price $>=p$.

BG (p): is the number of bids in H_M with price $>=p$

UAL (p): is the number of unaccepted asks in H_M with price $<=p$

Cubic-spline interpolation is used for prices for which no orders or trades are registered in H_M. The GD agent then chooses a price *p* that maximises its expected surplus, which is defined as the product of *f (p)* and the profit made by trading at price *p*.

GD agents are currently the best CDA agents as demonstrated by the results in Tesauro and Das (2001) which show that the GD agent strategy can dominate not only the ZI-C and ZIP strategies discussed but also the Kaplan agent strategy which had earlier been shown to dominate all agents submitted to a major comparative study in the Santa Fe Double Auction Tournament (SFDAT) which is described in Rust et al (1993). Moreover as we discuss in the next section GD agents have been shown to consistently outperform human traders by about 5-7% in laboratory experiments.

2.3 Agent-Human Interactions

Das et al (2001) describe experiments in which they allowed, for the first time, human subjects to interact with software bidding agents in a CDA. Their chief finding was that agents consistently obtain larger gains from trade than their human counterparts. Another significant and unexpected finding was the persistent far-from-equilibrium trading. This was in sharp contrast to previous work, including work by Smith (1962) Gode and Sunder (1993), Cliff (1997, 2001a, 2001) and several others in all-human or all-agent markets.

Das et al (2001) envision a future in which agents will play a dominating role in all trading marketplaces. Machine superiority in the CDA and other common auctions could have a direct and powerful financial impact – one that might be measured in billions of dollars annually. However for this vision to be realised: -

> "…it must be demonstrated that within their domain of application, agents can attain a level of economic performance that rivals or exceeds that of humans on average, without introducing undue risk" (Das et al, 2001, p. 1)

Das et al's experiments are aimed at comparing the performance level of current-day simple trading agents to human traders in realistic settings.

The experiments are conducted in a real time asynchronous CDA environment, which is provided by a special purpose distributed system called GEM that has been specially designed for conducting market experiments. The auctioneering mechanism has no way of distinguishing whether a message originates from a human or agent trader, both of which are given a uniform interface to the system. The ZIP and GD trading strategies (described in Sections 2.2.2 and 2.2.3 respectively) were used and minor modifications were made to both algorithms to adapt them to the auctioneering mechanism used. An equal number of agent and human traders were used in all experiments and the average surplus is evenly distributed between them: each human's set of limit prices is mirrored by one of the agents throughout any experiment. Thus, all attempts are made to ensure that neither side has any undue advantages. Humans are paid in proportion to their surplus to motivate better performance and the experiments are carried in 9 to 16 3-minute periods just as in previous work by Smith (1962), Gode and Sunder (1993) and Cliff (1997).

Das et al (2001) find that: -

1) There was significant interaction and trades between agents and humans, even though the agents were potentially faster. Roughly 30% of the trades happen between humans and agents. This is a reasonable percentage of the naive expectation of 50% if any trade partner (agent or human) is equally likely and the laboratory market does genuinely test human-agent interactivity as opposed to two non-interacting sub markets operating in different time scales.
2) When considered as a group, the agents outperform the humans in all experiments conducted. The total surplus obtained by agents extracted by agents is on an average 20% more than the surplus extracted by human traders.
3) Human performance does improve over the course of the experiment as the humans learn how to execute a good bidding strategy. Even then, there is a consistent edge in agent surplus over human surplus of at least 5-7% in the final periods of each experiment.

Das et al's (2001) work is a very strong demonstration supporting the claim that agents are, in general, capable of performing much better than human beings – all other things being equal. This has led Kephart (2002, p.1) to remark: -

> "Humans are on the verge of losing their status as the sole economic species on the planet. Software agents will [soon] represent – and be – consumer, producers and intermediaries."

If markets will be dominated by agents – is it then mandatory to use mechanisms that have specifically designed for human? Agents are willing participants in any market and we feel that in agent-dominated markets, there is an incentive to use novel mechanisms if they help us in achieving better performance on the whole. Not only can such new mechanisms be used in financial markets such as the CDA but also in applications of Market Based Control (MBC), which was discussed briefly in Section 2.1.3. We describe some work in this area in the next section. The thorough exploration of this idea forms the basis of the rest of our work.

2.4 Automating Market Design

To automate the process of market design we need to parameterise the description of a given market type. Once we have a complete parametric description of a market mechanism, the design process can be looked at as a parameter optimisation problem with the optimum parameter values specifying the best design.

Such an attempt was made by Cliff (2001a). Cliff had previously used a GA to optimise the parameters for his ZIP trading algorithm in (Cliff, 2001). A tuple of 8 parameter values defines a complete parameter set for initialising the ZIP trading strategy. Cliff uses a GA to optimise this tuple of 8 parameter values for traders participating in a CDA, thus automating the subjective task of selecting the right parameter values for initialising the algorithm.

In Cliff (2001a) he added a single parameter Qs to the tuple of ZIP initialising parameter values. This parameter Qs determines how an auction is carried out in the following manner. In a normal CDA, any buyer or seller can shout a bid or an offer at any given time. The way to implement this in a computer simulation of a CDA is to generate which trader will quote a price next uniformly at random from the pool of available traders. In a market in which the number of buyers and sellers is equal, this means that the probability of the next quote coming from a seller is 0.5. Similarly, the probability of the next shout coming from a buyer is 0.5 as well. If we let Qs denote the probability of the next shout coming from a seller, then for a normal CDA Qs=0.5. If we let Qb denote the probability that the next quote will come from a buyer, it is easy to see that Qb=1-Qs.

Thus, in terms of the framework we have described, Qs completely describes the market mechanism.

Since in an *English Auction* only buyers bid and sellers stay silent, in terms of this framework Qs=0. Similarly since in a *Dutch Auction* since only sellers shout offers, Qs=1. For an ordinary CDA Qs=0.5.[8]

But what do values like Qs=0.1 mean? The implementation of an auction mechanism in which Qs=0.1 is quite straightforward. In a mechanism with Qs=0.1 during the next 100 shouts, on an average only 10 will be from sellers and the remaining 90 from buyers. Perhaps, a more important question is whether there is any benefit in implementing such a curious mechanism. Cliff (2001b) seeks to answer precisely this question and reports that in 3 tests that he conducted on three different markets M1, M2 and M3, the optimum value of Qs that is evolved using a GA search are 0.0001, 0.07 and 0.16. Cliff (2001a, p. 1) comments on his results as follows: -

> "While there is nothing to prevent the GA from settling on solutions that correspond to the known CDA auction type [Qs=0.5] or the EA-like or DFA-like one-sided mechanisms, we repeatedly find that hybrid solutions are found to lead to the most desirable market dynamics. Although the hybrid 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."

Cliff also conducts trials with the value of Qs fixed at 0.5 to benchmark his results and reports results that confirm that the values evolved do indeed outperform an ordinary CDA market mechanism.

In a subsequent work, Cliff (2002) describes the results on test market in which the supply and demand schedule is changed suddenly between experiments. This is referred to as a *market shock* and seeks to mirror the occurrence in everyday financial market when a sudden event can lead to a massive change in trader preferences, thereby changing the underlying supply and demand curves significantly. In these experiments it is found that

---

[8] Since in an hybrid CDA any trader can accept a quote at any time, Qs=0 and Qs=1 are not strictly equivalent to the English or Dutch auctions. In the EA or DFA the silent party will wait until all bids or offers have been shouted and only accept the best one.

the optimum market mechanisms evolved are more strongly hybrid in nature (Optimum Qs values not close to either 1,0 or 0.5) than the results in Cliff (2001a). In the first experiment the first of the supply and demand schedules used in his earlier work (Cliff, 2001a) M1 and the second schedule M2 and switched midway – in this test M1M2 the optimum value of Qs settles to around 0.25. In the second experiment the reverse is done, schedule M2 is used to begin with and then switched with schedule M2. In this experiment the optimum Qs evolved is 0.45. For M1M2 a Qs value of 0.25 outperforms the standard CD variant with Qs=0.5. However, Cliff notes that for M2M1 the optimum value of Qs found by the GA (0.45) is no better (but also no worse) than the standard 0.5 variant.

However, there are several questions that could be raised about the validity of Cliff's results. Firstly, all experiments were performed with ZIP traders with the initialisation parameter values for the ZIP algorithm co-evolving alongside. Hence the exclusive use of the ZIP trading algorithm raises serious questions about the generalisation and robustness of these results to a change in the trading algorithm used.

Secondly, Cliff uses Smith's 'coefficient of convergence' α (discussed on Page 13) to measure the performance of the market mechanisms being tested. As has been discussed earlier in the context of ZI-C traders, high performance on one metric does not guarantee performance on another. Do Cliff's results generalise well to other fitness measures?

We seek to answer both these questions in our work. We have used ZI-C traders in our work, which are not only a different trading-algorithm from ZIP but also define the lower-bound on rationality required to participate in the CDA. We have already discussed some this and other benefits of using ZI-C traders in Section 2.2.1. We discuss our rationale for choosing ZIC traders in detail when we describe our experimental set-up in Section 3.1.

Secondly, we have used allocative efficiency as a measure of market fitness. This for two reasons. The first reason is simply that ZI-C traders do not always equilibrate and it is senseless to use convergence as a measure of fitness for a market populated with ZI-Cs. Gode and Sunder (1993) have shown that ZI-Cs perform poorly on a profit-dispersion measure as well. However, it has been shown that ZI-C traders can extract a large percentage of available surplus (~97%) in a market. This makes allocative efficiency the natural choice for measuring ZI-C populated marketplace performance. The second reason is that we believe that allocative efficiency is a more direct measure of the motives behind market design. Faster and more stable convergence do not mean much to a trader or to a market organiser on its own. It must be accompanied by other measures like an increase in average profit (from the point of view of a trader) or an average increase in the total profit made in the market or the percentage of viable trades executed (from the point of view of a market organiser). Allocative efficiency directly reports on the latter (and profit dispersion on the former).

It is worth mentioning the work by Phelps et al (2002) that was reported very recently. In their preliminary report Phelps et al (2002) discuss how they have used a Genetic Programming (GP) approach to co-evolve agent parameters as well as a trade-settling formula. Usually trades are settled at either the price of the earlier shout or midway between the bid and offer prices. In their work Phelps et al try to evolve an arithmetic expression using GP that can be used as an alternative trade-settling formula instead. They have used agents which use a myopic reinforcement-learning based strategy. Initial results show that formulae evolved by GP can give better performance over markets which use a conventional approach.

2.5 Summary

Markets can be described and characterised in terms of their underlying supply and demand schedules. Given a stable underlying supply and demand, the transaction price in the market will converge to a clearing or *Equilibrium Price* ***Po***, at which the quantity supplied is equal to the quantity demanded. There are various ways of measuring market performance including Allocative Efficiency, Smith's coefficient of convergence α and Profit Dispersion.

Some well-known and prevalent types of markets are the *Dutch Auction, English Auction, Vickery Auction and the Continuous Double Auction.* By far the most popular form of the DA is the *Continuous Double Auction (CDA)* or *open-outcry* market. In a CDA buyers and sellers simultaneously and asynchronously announce bids and offers: and at any given time a buyer is free to accept to an offer made by any seller and vice-versa. The CDA is an extremely important but little-understood institution and game-theoretic analysis of the CDA is intractable for most cases. However, experimental methods can be used to gain insight into the properties of the CDA. Trading-agent design for the CDA is a well-established field of research today and recently Das et al (2001) have shown that trading agents can even outperform human-traders. This has lent credibility to the paradigm of studying markets using agent-based simulations.

Market-mechanism design has so far been a highly subjective skill, more of an art than a science that requires considerable skill and experience along with a high degree of domain knowledge. We can automate the process of market design by using a parameterised description of market-mechanisms and gauge their fitness by using computational simulations of markets populated with trading-agents. Qs is one such parameter which determines how an auction is carried out by regulating the bidding process. Qs defines the probability of the next shout in the auction coming from a seller. Previous work has shown that hybrid market with Qs values not equal to 1,0 or 0.5 can lead to the most favourable market dynamics. However, these results have exclusively used the ZIP trading algorithm and Smith's coefficient of convergence α as the fitness measure. Our work seeks to test if hybrid markets can be shown to dominate more conventional variants robustly in markets populated with traders other than ZI-C and with different fitness measures.

Chapter 3

Evolutionary Design Optimisation

In this chapter we describe our experiments in which we have used a Self-adapting Evolutionary-Strategy (ES) based approach to optimise market-mechanism design for a market populated with ZI-C agents.
We present the first-ever results that demonstrate that hybrid and asymmetric market-mechanisms, in general, outperform standard market-mechanisms regardless of the type of trading-agents used.
We begin by describing our experimental set-up in Section 3.1. We explain the rationale for choosing ZI-C agents and describe how our methodology relates to that used in previous work including Smith (1962), Gode and Sunder (1993) and Cliff (1997).
In Section 3.2 we describe our algorithm in detail and discuss some of its parameters. We present the results of our experiments using this algorithm in our experimental set-up in Section 3.3. We describe the test cases, nature of fitness landscapes for these cases and detailed results for the behaviour of our EA on these cases.
Finally, in Section 3.5 we present the results of some rigorous tests that we have carried out to validate the results produced by the EA.

3.1 Experimental Set-up

We have exclusively used the ZI-C trader algorithm for all our experiments. Our methodology for organising a CDA is closely based upon the experimental set-up that was first used by Smith (1962) and subsequently followed by Gode and Sunder (1993), Cliff (1997) and others. It also incorporates some proposals made about the organisation of the CDA by Preist and van Tol (1998) and Tesauro and Das (2001).

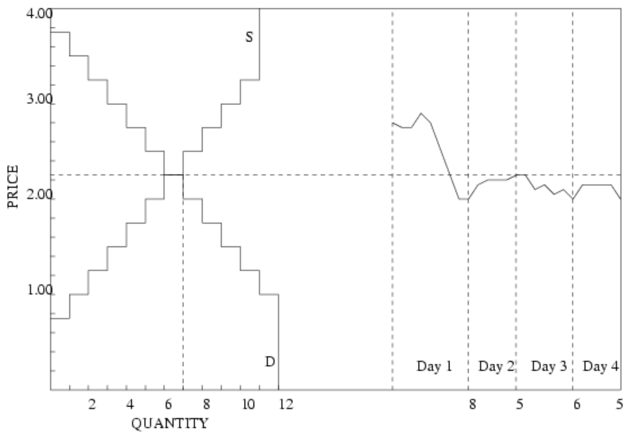

Fig 3.1 Smith's experiments were performed over many consecutive trading-periods or days. (Figure from Cliff (1997, Fig 9, p. 23). Originally from Smith (1962, Fig 8, p. 124))

3.1.1 Trading Days

Smith (1962) organised his experiments to be held over many consecutive trading periods or days. These varied from 2 to 6 days for different experiments. Trading units were replenished at the start of each trading period. The chief motive for this was to allow time for learning and a resulting convergence towards equilibrium price. Smith (1962, p. 114) states: -

> "…each trader's quotation is addressed to the entire trading group one quotation at a time…[over a period of days] there is an opportunity for traders to gain experience and to modify their subsequent behaviour in the light of such experience. It is only through some learning of this kind that I can imagine the possibility of equilibrium being approached in any real market."

Smith's results show a decreasing deviation of the transaction prices from P_0 in subsequent trading periods. Cliff (1997) also conduct their experiments over a number of consecutive trading days ranging from 10 to 30.

Cliff and Preist (2001) criticise the tradition of dividing trading in a CDA into periods and argue that in any actual CDA, trading cannot be divided into neat trading periods. They argue that this practice is unnecessary, but their results show that the division of time into trading periods does not have a significant effect on the stabilization of the marketplace.

Gode and Sunder (1993) use a fixed number of trading periods (6) for all their experiments but reported that even though there was a convergence to equilibrium during each trading day – there is no overall convergence on subsequent days. Since ZI-C traders do not learn: this is expected. Hence, it does not make sense to talk of subsequent trading days for ZI-C traders as their behaviour is guaranteed to be independent and uniform across all repeats of any experiment. Therefore we do not use the practice of dividing our experiments into trading days.

3.1.2 Result Validity

However, this random behaviour also raises a serious question about the validity of any simulation results obtained by using ZI-C traders populated markets. One way of reducing the variability in results obtained in ZI-C markets is to average results over a large number of repeated trials. Because of the highly random nature of ZI-Cs we are forced to perform many more repeats of any simulation that we conduct than would be required with any other 'intelligent' trading-agent algorithm. We typically average all results over thousands of repeats. We discuss this factor in more detail in Section 3.2.

To check any possible effects of the artefacts introduced by the random number generator algorithm used in our simulations on the results obtained we have tested our results with more than one type of random number generator. We have been able to successfully replicate our results with the standard C library random number generator provided by the GNU C compiler (Stallman, 2002), the Mersenne-Twister random number generator[9] (Matsumoto and Nishimura, 1998) and the random number generator described by Press et al (1992). However, we have generally used the standard C library random number generator for the sake of computational efficiency; the standard C library generator is faster than the other two by a factor of 10. Most results given in this report are generated using the standard C library random number generator.

3.1.3 ZI-C Traders

However the simplicity and random nature of the ZI-C trader algorithm is not without its advantages. The completely random behaviour of ZI-C traders is precisely the reason

[9] Code for the mersenne twister random number generator is available for download from the mersenne twister home page http://www.math.keio.ac.jp/matumoto/emt.html.

we have used them exclusively for all our experiments. Using ZI-Cs as opposed to any other trading-agent algorithm has two benefits.

Firstly, experimental results obtained with ZI-C trader experiments allow us to make a very strong statement about result-validity and robustness. Since the bids and offer space of ZI-C traders is a superset of all possible trading-strategies in any given market - they give us strong reason to believe that results obtained by using ZI-Cs are valid regardless of the trading-strategy used. In other words, the results are robust to changes in trading-agent strategy. Phelps et al. (2002, p. 3) discuss the importance of designing market-mechanisms that are robust to changes in trader strategy or incentives by drawing an analogy from co-evolution in nature: -

> "We view mechanisms as 'hosts' and the trading strategies as 'parasites'…it would be hoped that the mechanism population will adapt defences, and strategy-proof, incentive-compatible mechanisms would evolve"

Secondly, the original motivation behind the design on the ZI-C agent strategy was to address the perfect-rationality of trader-behaviour assumption. Since ZI-Cs represent the lowest end of the spectrum in terms of trader-rationality: it follows that any results obtained in ZI-C markets are robust to trader-irrationality. This is a fundamental concern in the field of market-mechanism design, as highlighted by Nicolaisen et al (2001). Market-mechanisms that perform well in spite of trader irrationality are naturally preferred over those mechanisms that deliver high-performance only with perfectly rational traders.

3.1.4 Number of Traders

Smith's original experiments were done with a varying number of traders with the numbers ranging from 11 to 24. This addressed an assumption made by most traditional economic models that a large (infinite) number of traders was needed for trends like equilibrium formation to emerge. Smith's results proved that a relatively small number of traders could approach and stabilize at equilibrium. Gode and Sunder (1993) and Cliff (1997) also use a small number of traders: Gode and Sunder use 12 and Cliff 22 traders. We show in Section 3.2 that the expression for our simulation algorithm's computational complexity contains a factor equal to the number of trading agents. So an experiment with 10 traders takes double the time taken by an experiment with 5 traders. Hence, for the sake of computational efficiency we limit the number of trading agents in all our test cases to less than or equal to 22.

3.1.5 Persistent Shout Double Auction

In implementing the actual bidding process we have incorporated two suggestions made by Preist and van Tol (1998) and followed by subsequently followed by Tesauro and Das (2001) and Das et al (2001) in their experiments. Preist and van Tol (1998) studied a variant of the CDA called the *persistent shout double auction market* in which a trader's current bid or offer persists until the trader makes another or is able to execute a trade at that price. Preist and van Tol (1998) demonstrate that agents in a persistent shout CDA reach equilibrium much faster, maintain a more stable equilibrium and are more robust to changes in learning rate (for ZIP agents).

The New York Stock exchange is a persistent shout CDA with an additional constraint that any new bids or offers must improve on the existing ones. This constraint is commonly referred to as the *NYSE rule* in trading-agent literature.

Thus, we implement a market which is a persistent shout CDA organised as an order queue (Preist and van Tol, 1998, p. 6). We have also incorporated the NYSE rule into our trading mechanism. This means that the best outstanding bid (highest bid) or offer (lowest offer) is either maintained or bettered till a trade occurs. We maintain bid and offer lists as priority queues with the bid queue arranged in descending and the ask queue in ascending order. The tops of both these queues are polled after every bid or offer to check if there are any outstanding trades waiting to happen. All outstanding trades are settled after every bid or offer. The transaction price is fixed at the middle of the bid and ask prices. In previous experiments transaction prices are either fixed at the earlier of the bid or offer, or fixed midway between the two. Phelps et al (2002) have recently proposed co-evolving a formula for settling the transaction price whenever a bid and offer cross.

The trading-periods first used by Smith (1962, p. 112) defined a finite period of time fixed in advance at the end of which all trading was stopped and un-traded units were declared void. Gode and Sunder (1993, p. 122) also use absolute time periods for their experiments but these time limits are different for human and software-traders (Humans were given 4 minutes and software-traders were given 30 seconds). This mechanism raises questions about the validity of any comparisons drawn between human and software trader performance, as in any real marketplace one of the major advantages that software traders have is that of speed of execution – which is many orders of magnitude faster than their human counterparts. We feel that number of shouts allowed is a better way of controlling the period of trading in an agent market. We use a parameter that fixes the maximum number of shouts allowed for a given market as an expression of the total number of traders in the market. Typically we use a factor of 100 (so for Qs=0.5 on an average each trader would get 100 shouts) but this is a parameter (MAX_SHOUTS) that can be configured at compile time. Cliff (1997) use an alternate mechanism that is motivated by the desire to execute all possible trades in a given market. Trading continues till either no feasible trades remain in the market or is there have been a fixed number (100) of continually rejected bids or offers.

3.2 Description of Evolutionary Algorithm Used

We have used a Self-Adapting Evolutionary Strategy (ES) style approach to design an Evolutionary Algorithm to optimise the value of Qs for a given market. Evolutionary Strategies are the application of the Evolutionary Computation (EC) paradigm to real valued parameter optimisation (Schwefel 1981, Schwefel 1995). Even though there are considerable overlaps between techniques in the Evolutionary Computation paradigm, including Genetic Algorithms (GA) (Holland, 1992 and Goldberg, 1989), Evolutionary Programming (EP) (Fogel et al 1966 and Fogel 1991) and Evolutionary Strategies (ES), the type of representation, selection schemes and search operators define the type of EA. ES are marked by the use of real-valued representation and the exclusive use of mutation as a search operator.

ES are defined by the population size, μ, and number of offspring generated in each generation, λ and the selection mechanism used. If the best μ individuals are now chosen from the pool of the parent (μ) and offspring (λ) to form the next generation of μ individuals the ES is known as a (μ+λ) ES. If the best μ individuals are now chosen only from the offspring (λ) population (here $\lambda >= \mu$) the ES is known as a (μ, λ) ES.

According to this characterisation our EA can be defined as a $(\mu+\mu)$ ES. The selection mechanism used, as in all ES, is a strictly rank-based elitist selection mechanism.

ES were designed to be self-adaptive from inception and studies by Fogel (1992, 1995) and Back and Schwefel (1993) have shown that an ES with a self-adaptive mutation operator usually perform better than an ES without a self-adaptive operator on the same problems. A general self-adaptive strategy, as outlined by Yao et al. (1999), is shown in Fig. 3.2.

1. Generate a random initial population of individuals. Each individual is taken as a pair of real-values $(x_i, \eta_i)\ \forall i \in \{1,2\ldots\mu\}$, where **x** is the value being optimised and η is the standard deviation for Gaussian mutations. η is also known as the **strategy paramater.**
2. Evaluate the fitness for each individual $(x_i, \eta_i)\ \forall i \in \{1,2\ldots\mu\}$ of the population based on the objective function $f(x)$.
3. For each parent (x_i, η_i) generate an offspring (x_i', η_i') as follows: -

   $$x_i' = x_i + \eta_I N(0,1)$$

   $$\eta_i' = \eta_I\, e^{(\tau' N(0,1) + \tau N(0,1))}$$

   where $N(0,1)$ is a normally distributed random number with mean 0 and standard deviation 1. The factors τ and τ are commonly set to $(2(n)^{1/2})^{-1/2}$ and $(2(n))^{-1/2}$.
4. Calculate fitness for each offspring (x_i', η_i')
5. Select μ individuals from the union of parents $(x_i, \eta_i)\ \forall i \in \{1,2\ldots\mu\}$ and offspring $(x_i', \eta_i')\ \forall i \in \{1,2\ldots\mu\}$.
6. Stop if halting criterion is satisfied else go back to step 3.

Fig 3.2 Schematic description of a typical Self-Adaptive Evolutionary Strategy (ES) or Evolutionary Programming (EP) style approach

The idea is to take longer steps across the fitness landscape in the beginning when the population is far away from the global optimum but take decreasingly small steps towards the end of the ES to hone onto the exact global optimum. A self-adaptive ES automates this process and lets evolution decide what step sizes are appropriate.

For example given 2 individuals (x_1, η_1) and (x_2, η_2) where $x_1=x_2$: if both individuals are currently far away from the global optimum, the chances of producing fitter offspring are greater for the individual with the larger step size η_i.. Hence, in the beginning when most solutions are far away from the global optimum larger step sizes will dominate. The opposite is true when the population has reached close to the global optimum and very fine steps are needed to hone into the exact optimum. Here larger step sizes will only take an individual away from the optimum – hence smaller values of the step size, η, will be advantageous. We show some plots of η values for successful offspring in Section 3.3.3 and demonstrate that in general step sizes decrease as evolution proceeds.

Figure 3.3 shows pseudo-code for the evolutionary algorithm that we have designed to optimise Qs values for a given market.

1) Initialise population **Qs[μ]** with *U[0,1]* where *U[0,1]* is a uniformly distributed random number between 0 and 1.
2) Evaluate fitness for all members and store in **FitnessList[μ]**
3) For **NO_OF_GENERATIONS** gens do
   a) For each member of population **Qs[i]**
      i) Produce corresponding Offspring **Qs'[i]** by mutating **Qs[i]**
      ii) Evaluate Fitness for all **Qs'[i]** and store in **OffspringFitnessList[μ]**
   b) Select the best μ members from **Qs'[μ]** and **Qs[μ]** and replace worse members in **Qs[μ]** with new better members from **Qs'[μ]** (Prefer member from Offspring population is the fitness is equal). Copy corresponding fitness values from **OffspringFitnessList[μ]** to **FitnessList[μ].**
   c) Log best fitness value in population **Qs[μ]** in fitness log file. Also log Qs value corresponding to best fitness value in population. Log η values for all successful mutations.

Fig 3.3. Pseudo-code for our Evolutionary Algorithm for optimising Qs value for a given market

Fitness is measured by conducting a simulation of a CDA in the given market with the value of Qs for which fitness is being evaluated and measuring the allocative efficiency at the end of the auction. It is generally a good idea to average measurements with an element of stochasticity over a number of trials. Since the behaviour of ZI-C agents is highly unpredictable we must average fitness measurements over a large number of trials. The number of these trials is a parameter of the algorithm, REPEAT_TRIALS and influences the accuracy of the results quite strongly. We discuss the effect of this parameter in Section 3.3.2.

3.2.1 Algorithm Complexity

It can be shown that the time complexity of the Evolutionary Algorithm shown in Fig 3.3 is **O (gprmb)** where: -

- **g** or MAX_GENS is the number of generations

- **p** is the population size
- **r** or REPEAT_TRIALS is the number of times a market simulation is repeated for the results to be averaged to give the fitness value
- **m** or MAX_SHOUTS is the maximum number of shouts allowed to each buyer on an average for Qs=0.5. In other words the total number of shouts is = (number of buyers + number of sellers) * m
- **b** is the total number of buyers/sellers in the market. The above expression assumes that the number of buyers and sellers in a market are equal (this is the case for all our test cases). [10]

We point out both the values for each run of the EA when we discuss our results in Section 3.3. We also discuss how changing each parameter affects the performance of the EA and the accuracy of the results. Of special concern are, **g**, the number of generations we allow our EA to run for and most importantly **r**, the number of repeat trials we use to strike a balance between algorithm performance and result accuracy. The importance of this parameter will become clear in Section 3.3.2 where we describe how a very high number of repeats leads to more accurate results, but can also be very expensive computationally, with the help of fitness landscapes.

3.3 Results

We now report results on 6 sample markets on which we have conducted experiments. Three of these – M1, M2 and M3 have been used previously by Cliff (2001a) in his work and we compare our results with those obtained by him. Three new markets – SDI, SDII and SDIII are also used

Perhaps the most significant finding of our experiments is that in most markets the value of Qs found to be optimum using an ES indicates that are strongly hybrid in nature i.e. they deviate from the conventional variants of 1, 0 and 0.5 significantly. We further confirm this by plotting fitness landscapes for all test cases and confirm the validity of the results given by an evolutionary search through the space of all possible market-types.

---

[10] In a more general case where the number of buyers and sellers are not equal we can rewrite the expression for time complexity as **O (gprmt)** where t is the greater of the number of buyers or sellers.

Parameter Settings

All results are computed with the following parameter set: population size 30, every fitness evaluation averaged over 1000 repeats and each ES is left running for 500 generations.

We show logs of Qs, fitness and step size (η) for each test case. We have repeated each experiment for test cases SDI to SDIII 25 times and experiments for test cases M1 to M3 10 times and we report the averages for each of these values (with plus and minus one standard deviation) as the evolution proceeds.

3.3.1 Description of test cases

The buyer and seller reserve prices for each of our test cases M1 to M3 and SDI to SDIII are shown below.

The test cases SDI to SDIII have 6 buyers and sellers each. The limit prices for buyers and sellers in each of these markets are given below.

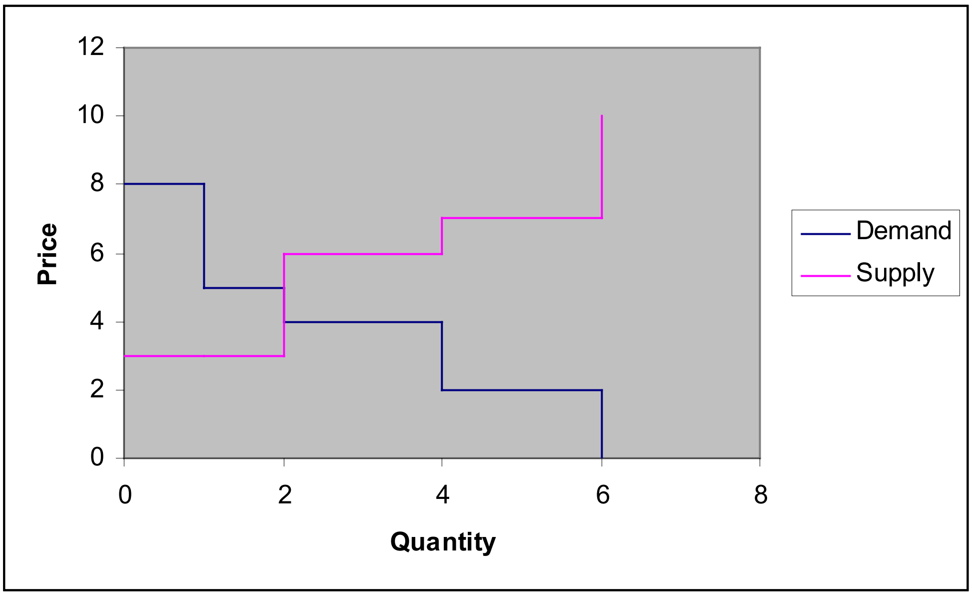

Buyers	8	5	4	4	2	2
Sellers	3	3	6	6	7	7

Fig 3.4 The reserve prices for buyers and sellers in SDI and the corresponding supply and demand curves. In this market Po=4 and Qo=2

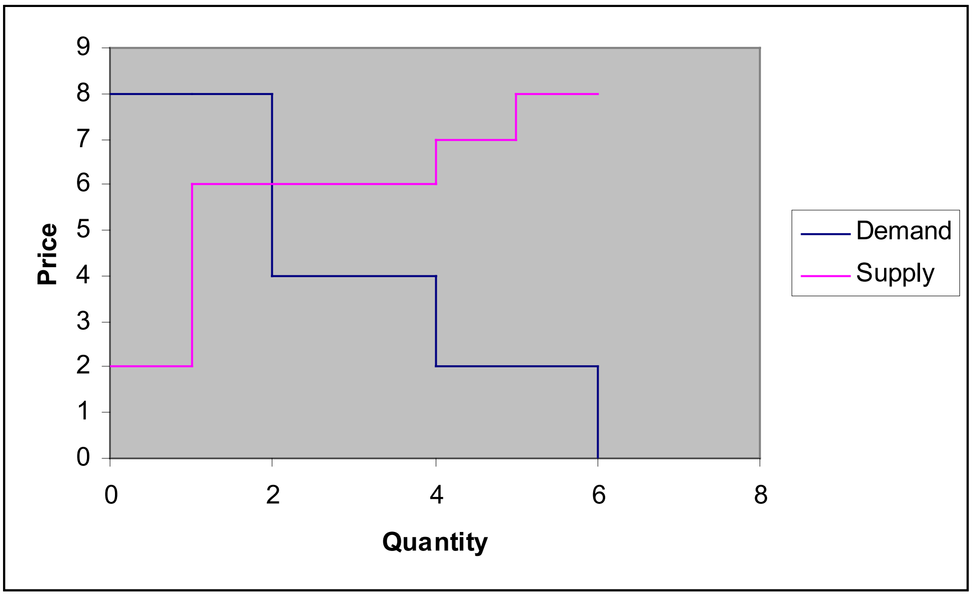

Buyers	8	8	4	4	2	2
Sellers	2	6	6	6	7	8

Fig. 3.5 The reserve prices for buyers and sellers in SDII and the corresponding supply and demand curves. In this market Po=7 and Qo=2

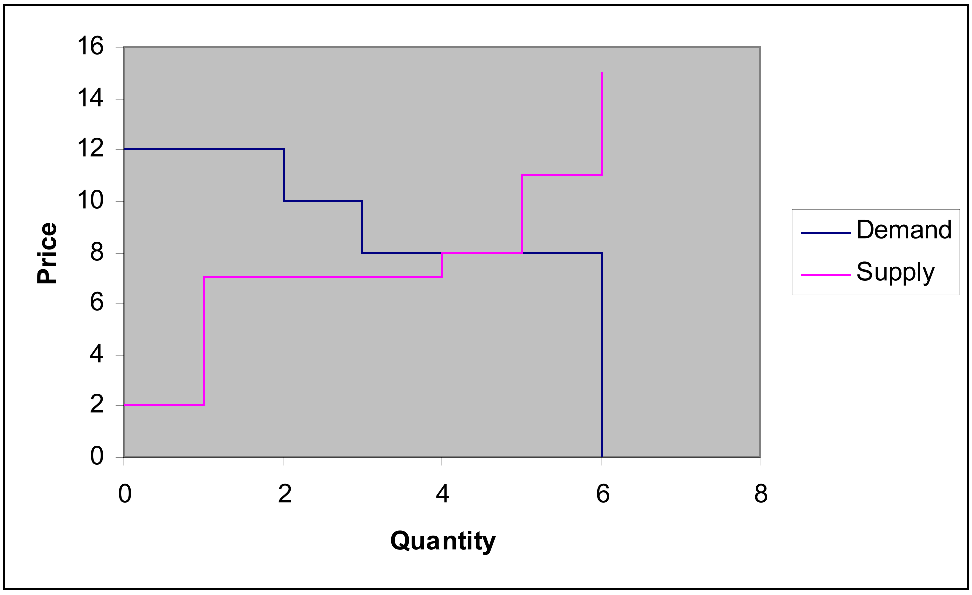

Buyers	12	12	10	8	8	8
Sellers	2	7	7	7	8	11

Fig 3.6 The reserve prices for buyers and sellers in SDII and the corresponding supply and demand curves. In this market Po=8 and Qo=5

The test cases, M1 to M3, were used by Cliff in his work on automatic market-mechanism design (Cliff, 2001a). The supply and demand curves for these 3 markets, as well as the values for the buyer and seller reserve prices are shown below.

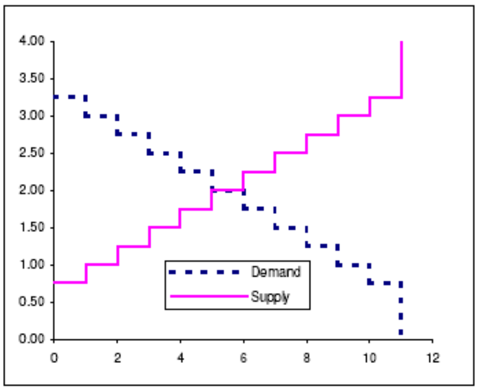

Figure 3.7 The supply and demand schedule for test case **M1**. The market consists of 11 buyers and 11 sellers. The equilibrium price Po = 2 and the equilibrium quantity Qo = 6

Buyers	3.25	3	2.75	2.5	2.25	2	1.75	1.5	1.25	1	0.75
Sellers	0.75	1	1.25	1.5	1.75	2	2.25	2.5	2.75	3	3.25

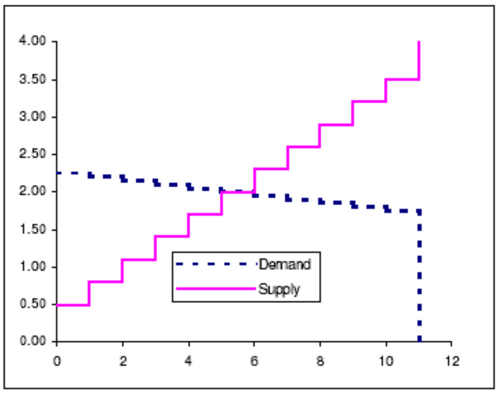

Figure 3.8 The supply and demand schedule for test case **M2**. The market consists of 11 buyers and 11 sellers. The equilibrium price Po = 2 and the equilibrium quantity Qo = 6

Buyers	2.25	2.2	2.15	2.1	2.05	2	1.95	1.9	1.85	1.8	1.75
Sellers	0.5	0.8	1.1	1.4	1.7	2	2.3	2.6	2.9	3.2	3.5

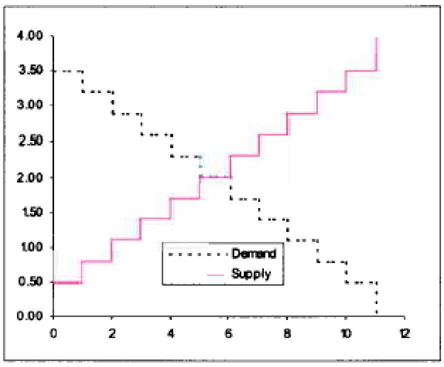

Figure 3.8 The supply and demand schedule for test case **M3**. The market consists of 11 buyers and 11 sellers. The equilibrium price Po = 2 and the equilibrium quantity Qo = 6

Buyers	3.5	3.2	2.9	2.6	2.3	2	1.7	1.4	1.1	0.8	0.5
Sellers	0.5	0.8	1.1	1.4	1.7	2	2.3	2.6	2.9	3.2	3.5

We hope to convince the reader of 2 main points in this section.

First, that the ES we have designed can find the correct optimum Qs value for all the test cases. We do this primarily by checking whether the optimum Qs value found by the ES lies on the peak of the underlying trend that emerges from the landscapes we have plotted for each problem.

Second, that most of the markets we have studied are hybrid variants of the CDA with Qs not equal to either 1, 0 or 0.5 offer the largest expected gains in them.

3.3.2 Fitness Landscapes

Fitness landscapes can be generated quite easily for our test problems by stepping through a series of Qs values from 0 to 1 and evaluating fitness for these values. Even though these landscapes can be prohibitively expensive to generate, especially for large values of REPEAT_TRIALS, they are extremely useful for validating EA results and evaluating the behaviour of the EA on test problems.

Figure 3.10 shows the fitness landscapes for problem SDI for different values of the REPEAT_TRIALS parameter. We briefly discussed the effect of this parameter on the time complexity of our EA in Section 3.2.1. In the plots shown below we have plotted the fitness values for a 1000 values of Qs between 0 and 1. Figure 3.10 clearly shows the effect this parameter has on the accuracy of the underlying fitness landscape. Increasing the number of trials leads to an increasingly accurate estimate of the true underlying landscape.

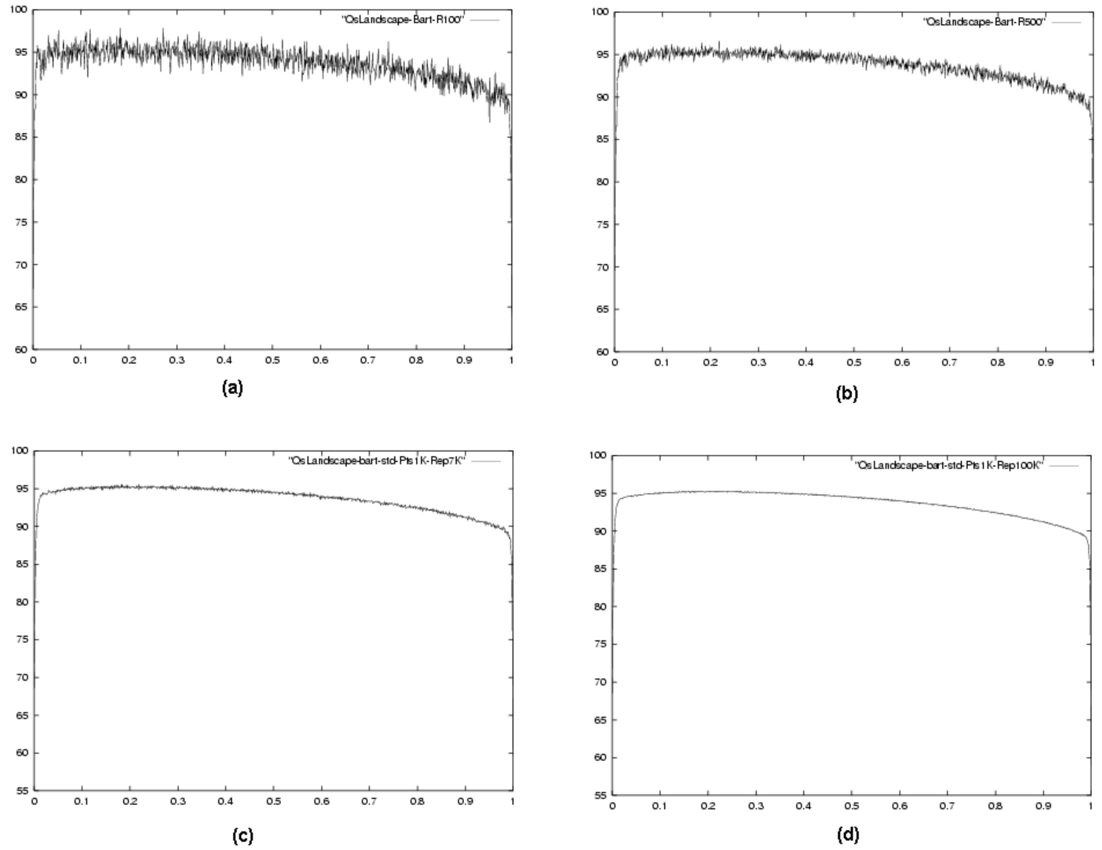

Fig 3.10. This figure shows the fitness landscape for the SD-I problem for different values of REPEAT_TRIALS. The fitness value is averaged over a large number of trials and an increase in this number leads to an increasingly accurate estimate of the true underlying fitness landscape. Part (a) shows the landscape for SD1 in which fitness has been averaged over a 100 trials. Similarly parts (b), (c) and (d) show landscapes for 500, 7,000 and 70,000 trials.

We can see from 3.10 (a) that there is a strong underlying trend in the fitness landscape for SD1. Increasing the number of repetitions leads to confirmation of this with Figures (b), (c) and (d) depicting a landscape with an increasingly clear underlying trend.[11]

Time Complexity for Generating Landscapes

It can be shown that the time complexity for generating fitness landscapes for markets is **O (p r m b)** where **p** is the number of Qs points we evaluate fitness for and **r**, **m** and **b** have the same meaning as described Section 3.2 on page 43.

---

[11] Note that generating a fitness landscape for SD1 for 1000 values of Qs with a value of REPEAT_TRIALS=1000 takes approximately 12 hours on an Intel Pentium-4 1.8 GHZ PC with 256 MB RAM.

It should be clear that an approach that generates highly accurate fitness landscapes to then use a hill-climbing algorithm to find the optimum value of Qs for a given market would be prohibitively expensive. Figure 3.11 shows that even landscapes with a high-degree of accuracy do not lend themselves to this approach.[12]

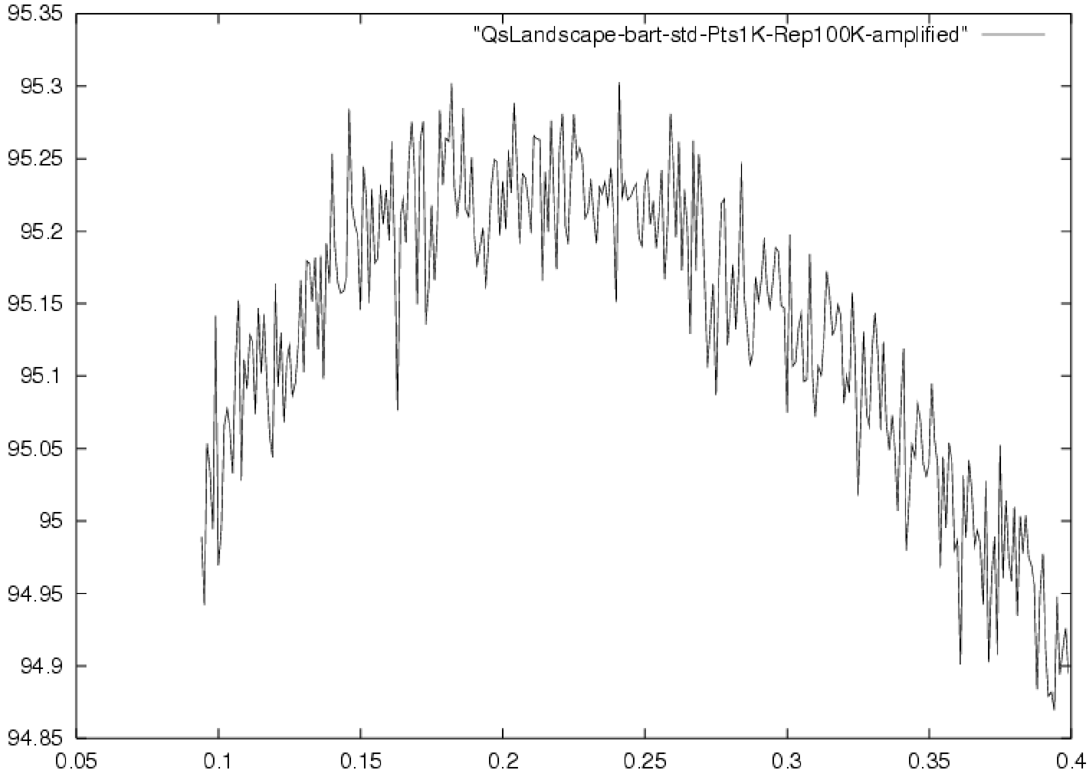

Fig 3.11 A magnified view of Figure 3.10(d). This graph only plots values for Qs ∈ [0.1,0.4] and demonstrates that even if we generate fitness landscapes with a high-degree of accuracy the underlying trend is still too noisy for applying a straightforward hill-climbing approach to it.

[12] This landscape took approximately 10 days to compute.

3.3.3 SD I

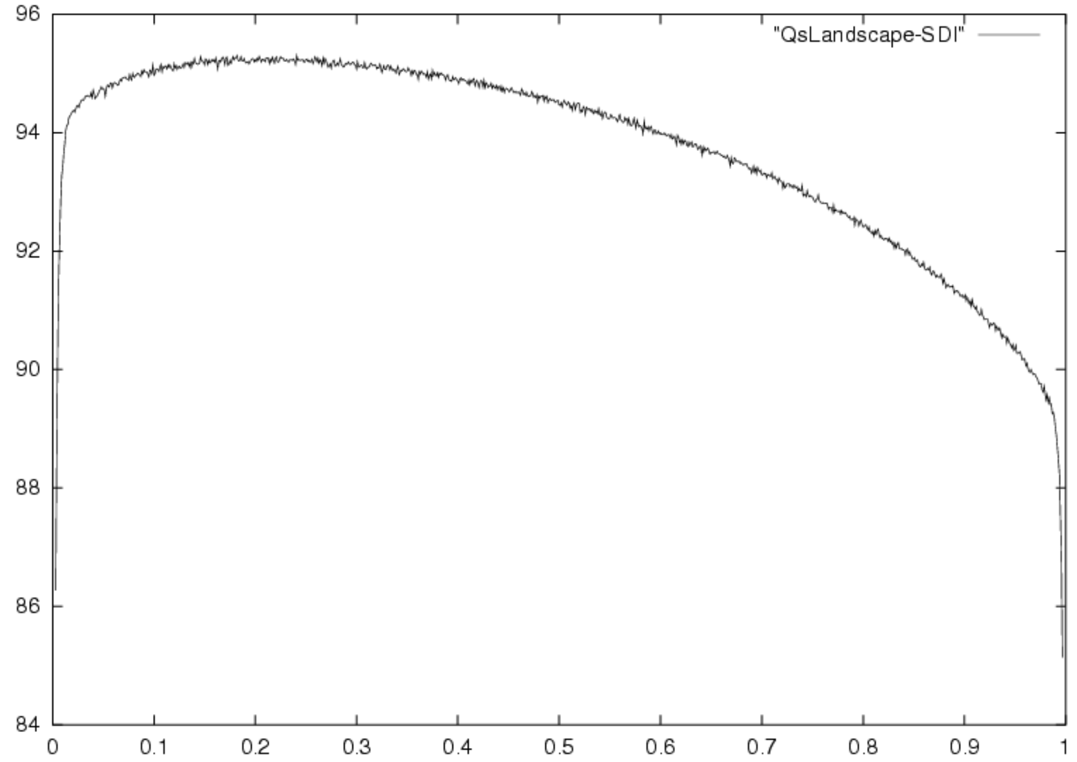

Figure 3.12 The fitness landscape for the SDI problem. Note that this landscape has been rescaled and plotted for fitness values above 85%. The fitness values are averaged over 10,000 repeats.

Even by visual inspection one can observe that the SDI landscape is strongly asymmetric and the optimum lies at none of the conventional values of Qs equal to 1, 0 or 0.5.

The step size (η) log shown in Figure 3.13 confirms the predicted behaviour of the self-adaptive ES as the step sizes for successful offspring get smaller and smaller as the ES hones into the exact optimum. Note that the step size η does not give the value of the actual mutation. The actual mutation value is given by $\eta * N(0,1)$, where $N(0,1)$ is a normally distributed random number with mean 0 and standard deviation 1.

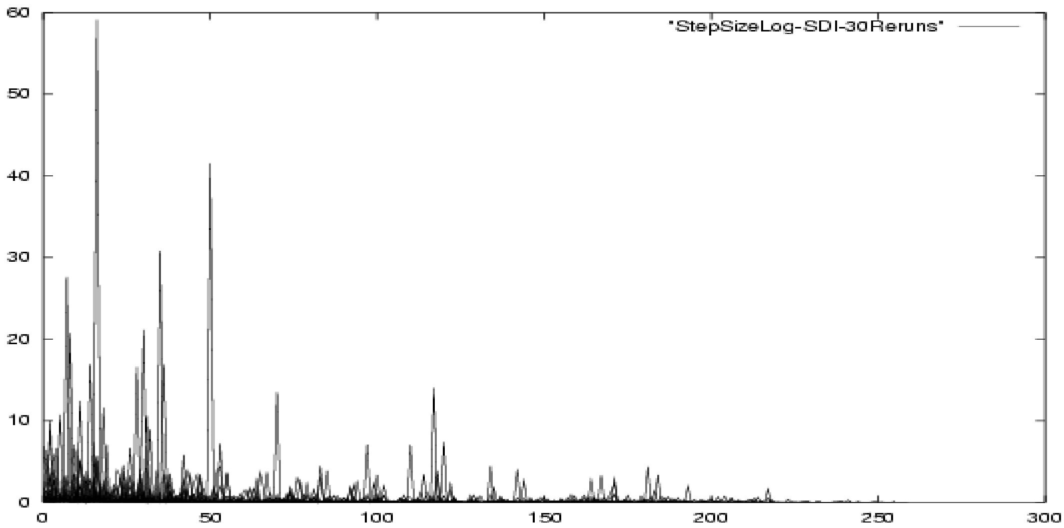

Fig 3.13. This figure shows the log of the Step Size (η) values for successful offspring as evolution proceeds for all 25 runs of the ES on test case SDI.

The log for elite Qs values shown in Figure 3.14 shows the optimum to lies around Qs=0.2 (This log depicts information averaged over 25 runs). This agrees with the information depicted in the landscape for the problem, which was shown in Figure 3.12. Note that the standard deviation is around 0.05 in this case.

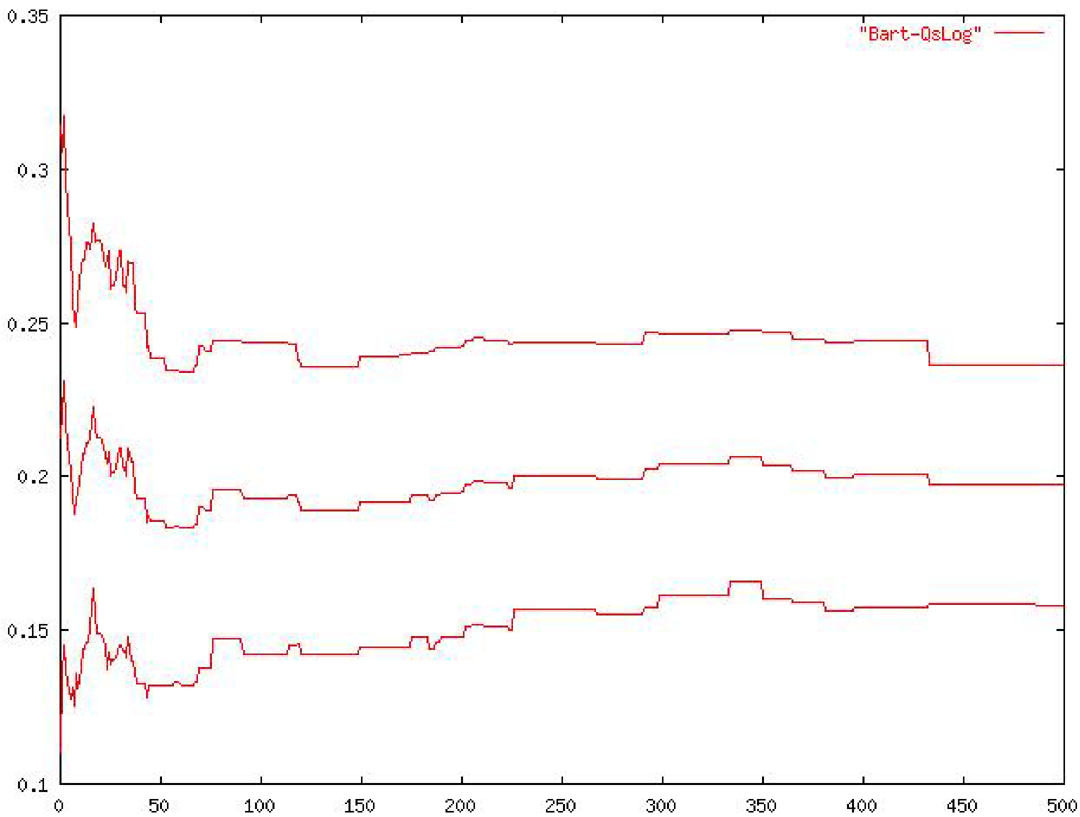

Fig 3.14 The log of elite Qs values averaged over 25 runs for SDI. The lines for plus and minus one standard deviation are also shown.

The corresponding log for elite fitness is shown in Figure 3.13.

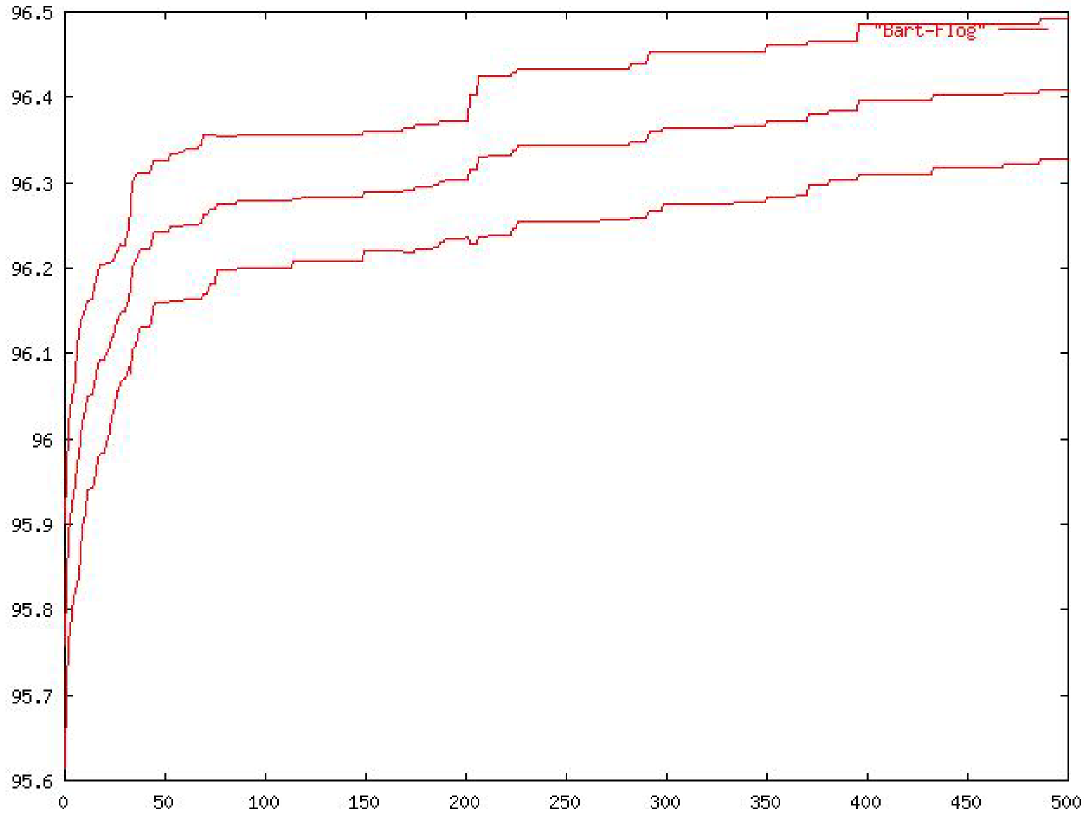

Fig 3.15 The log of elite fitness values averaged over 25 runs for SDI. The lines for plus and minus one standard deviation are also shown.

3.3.4 SD II

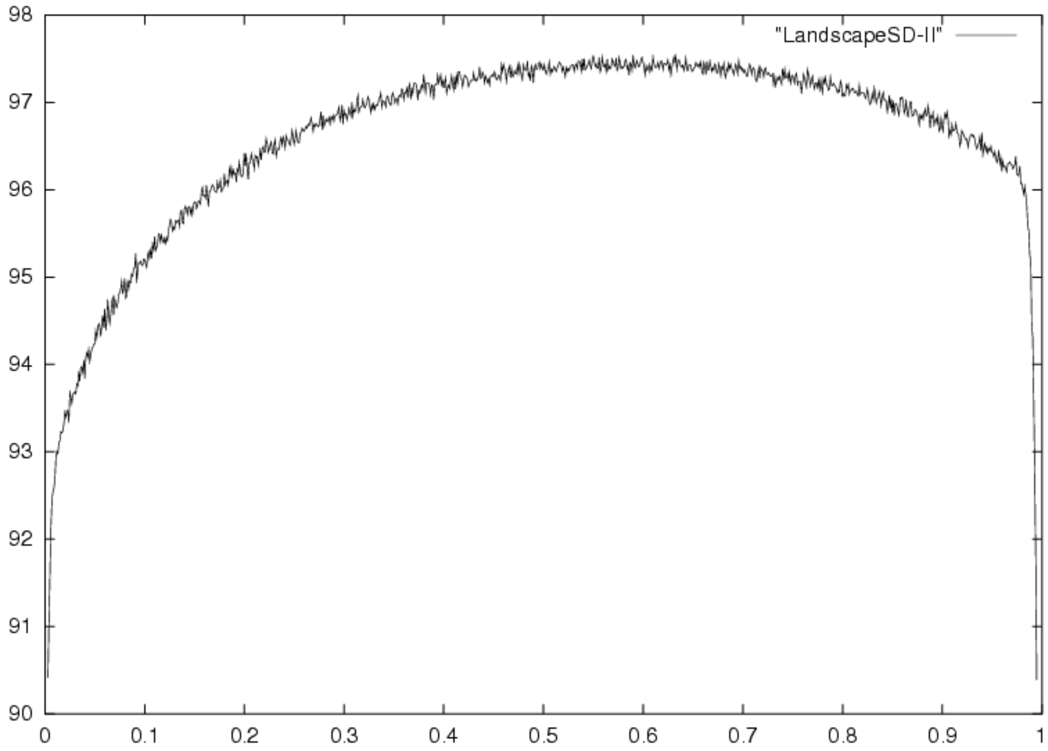

Fig 3.16 The fitness landscape for the SDII problem. Note that this landscape has been rescaled and plotted for fitness values above 90%. The fitness values are averaged over 10,000 repeats.

The landscape for SDII is shown in Figure 3.16. Once again the Qs log shown in Figure 3.17 agrees with the information depicted by the landscape and the optimum lies at about Qs=0.58. The standard deviation in the runs is about 0.04. The corresponding fitness log is shown in Figure 3.18

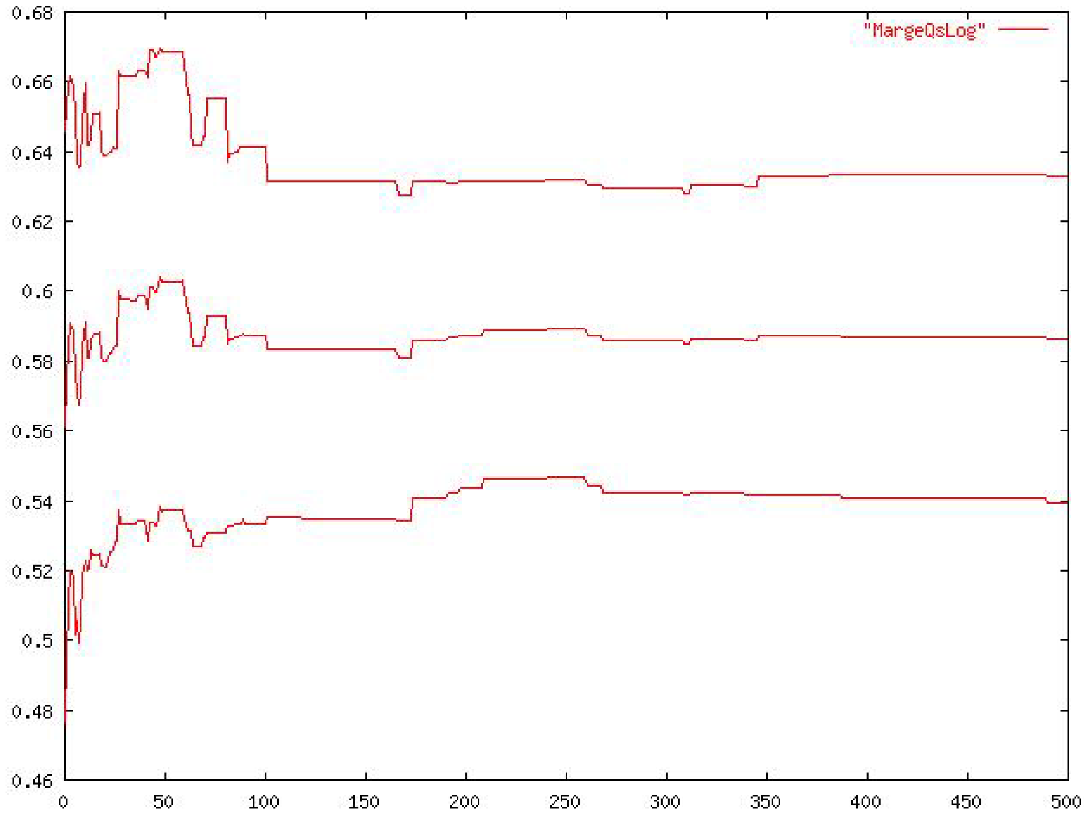

Fig 3.17 The log of elite Qs values averaged over 25 runs for SDII. The lines for plus and minus one standard deviation are also shown.

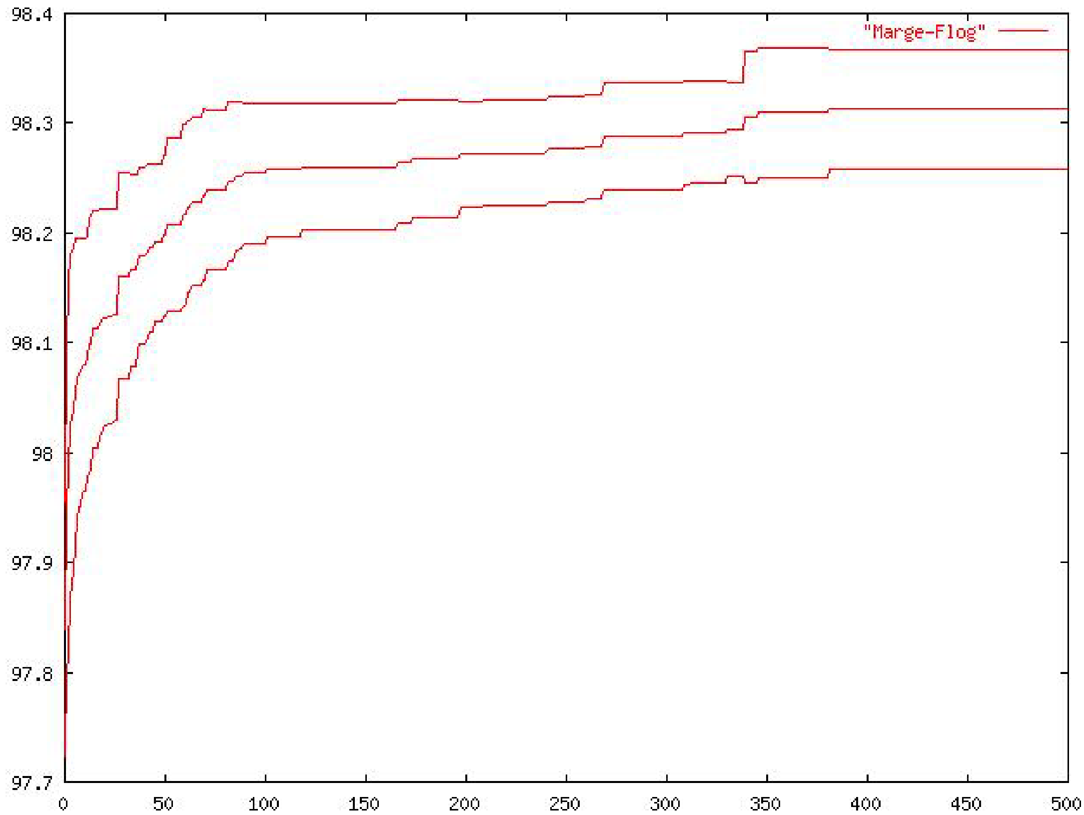

Fig 3.18 The log of elite Fitness values averaged over 25 runs for SDII. The lines for plus and minus one standard deviation are also shown.

3.3.5 SD III

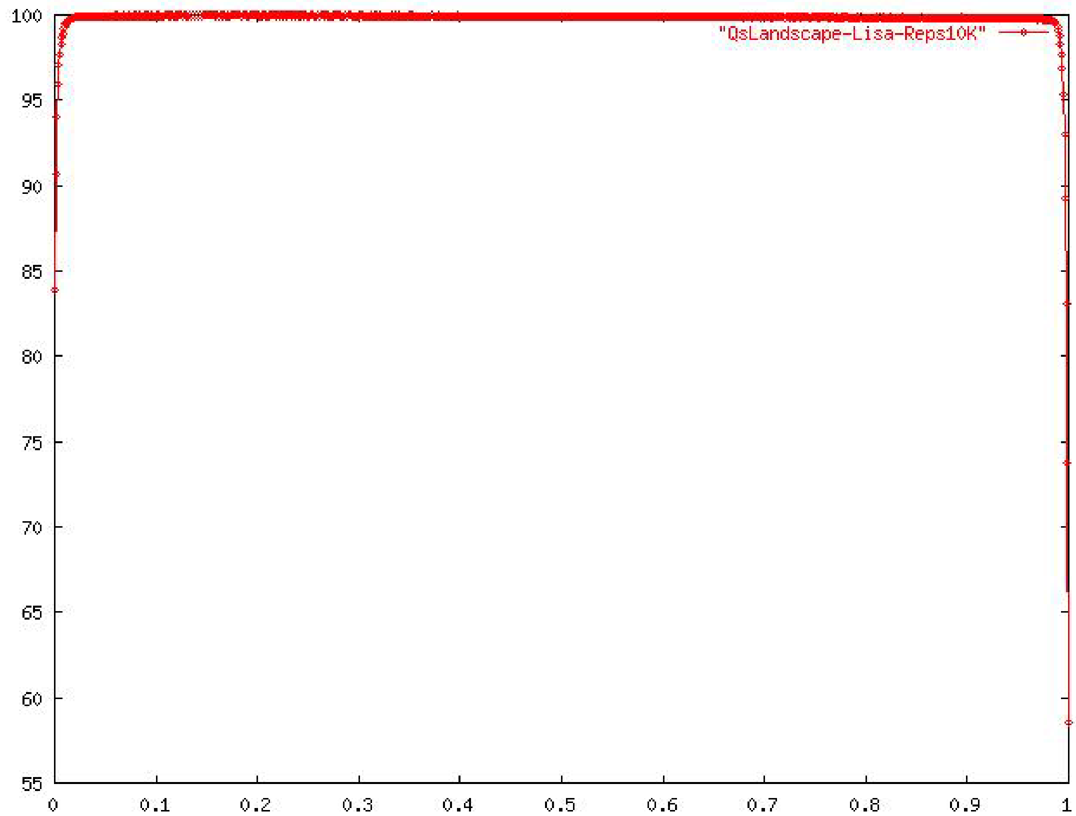

Fig 3.19 The fitness landscape for the SDIII problem. This landscape is conspicuous because of its seemingly linear nature with most points having 100% fitness. The fitness values are averaged over 10,000 repeats.

The landscape for SDIII shown in Figure 3.19 immediately stands out because of its seemingly linear nature and because of the fact that most of the points in the landscape seem to have a 100% efficiency. The elite Qs values for all 25 runs, which are shown in Figure 3.21, is extremely noisy and shows poor convergence due to this highly flat landscape, which behaves like a plateau over which the ES displays an unstable behaviour. However, note that the values are clustered around Qs=0.2.

Buyers	12	12	10	8	8	8
Sellers	2	7	7	7	8	11

Figure 3.20 Buyer and Seller limit prices for market SDIII

A qualitative explanation of why the base-line efficiency in this market is so high can be given as follows. Consider the supply and demand schedule for the SDIII problem reproduced in Figure 3.20. The only extra-marginal trader in this market who can reduce efficiency by executing a trade is the seller with reserve price 11. Furthermore, the only buyer he can trade with has a reserve price of 12. Given that this buyer always shouts prices below 12 (distributed uniformly in the [0,12] interval) and the extra-marginal seller always shouts prices above 11 – it is highly unlikely that they will ever trade. Hence, a 100% efficiency is achieved in most runs. A full quantitative analysis for calculating this likelihood is hard due the presence of the market queue, as well as the fact that we must consider the actions taken by the other traders. But the qualitative explanation gives us a fair understanding of the reason behind the nature of the landscape.

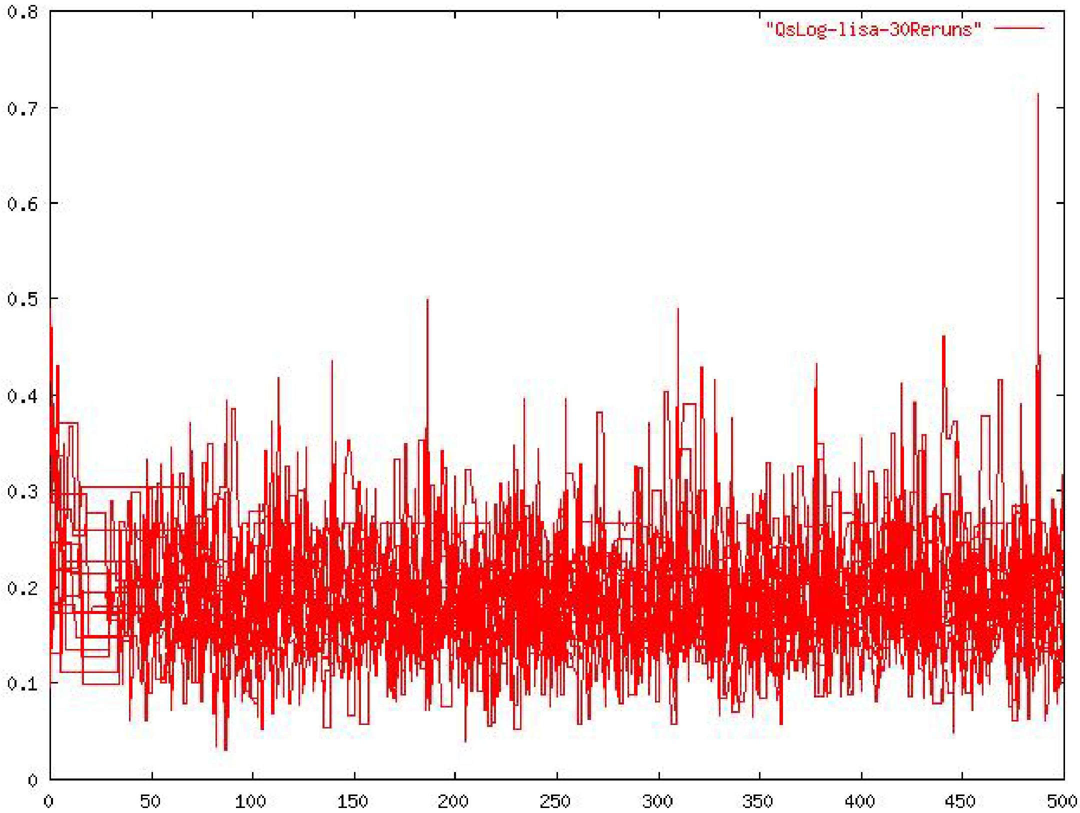

Figure 3.21 This figure shows the elite Qs value log for all 25 runs of the EA

The step size log shown in Figure 3.22 corroborates the hypothesis about the poor convergence of the ES on SDIII as the step size values do not show a steady decrease in magnitude.

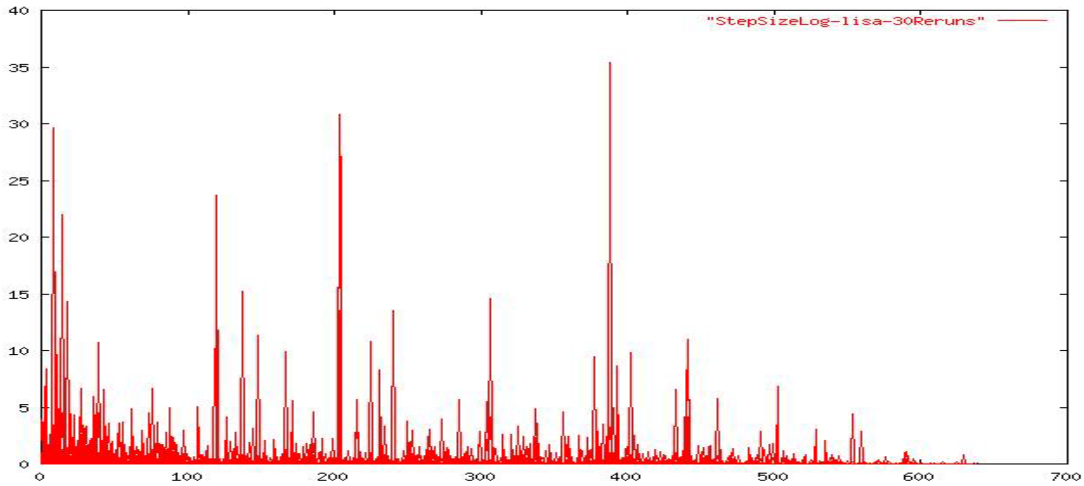

Figure 3.22 This figure shows the log of the Step Size (η) values for successful offspring as evolution proceeds for all 25 runs of the ES on test case SDIII

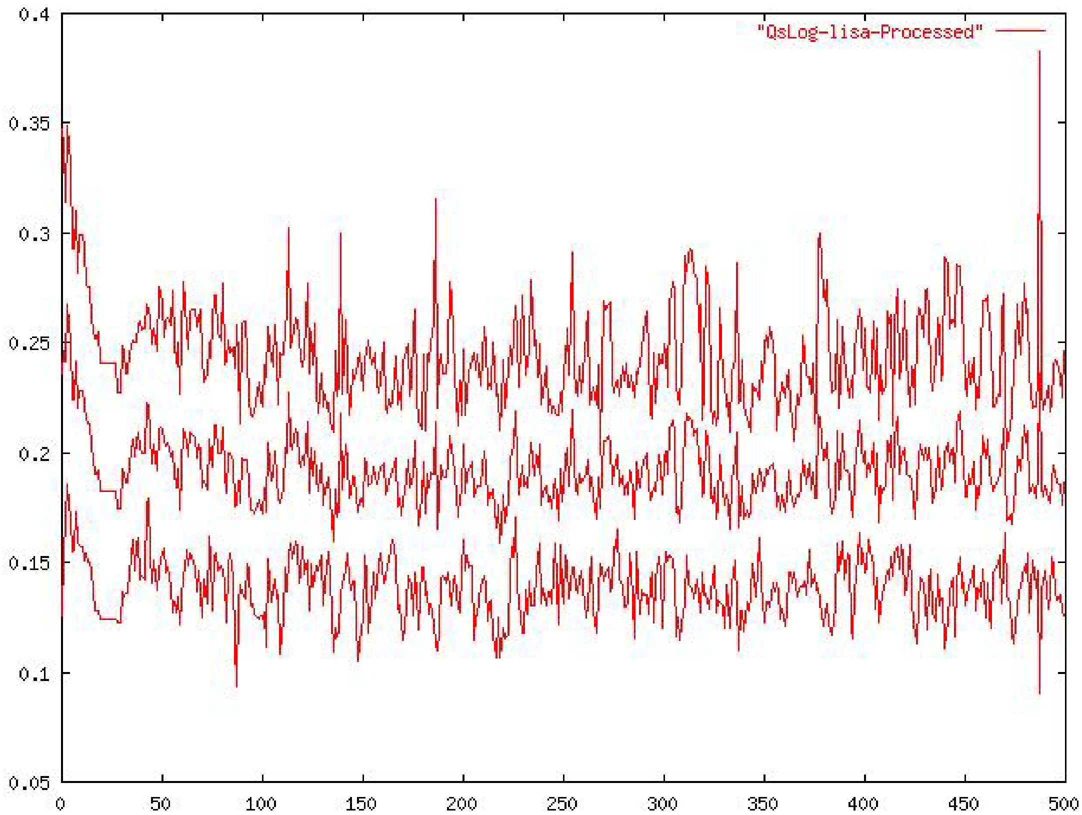

Figure 3.23. This figure shows the log for the elite Qs value averaged over 25 runs of the ES. The optimum lies around Qs=0.2. Note the noisy convergence of the ES.

The plot for the average elite Qs values is shown in Figure 3.23. The optimum seems to be around at Qs=0.2 but this cannot be confirmed by the extremely flat landscape shown in Figure 3.19. Moreover the noisy convergence of the ES raises even more doubts about the validity of this result. However the elite fitness log shown in Figure 3.24 confirms the information depicted in the landscape shown in Figure 3.19.

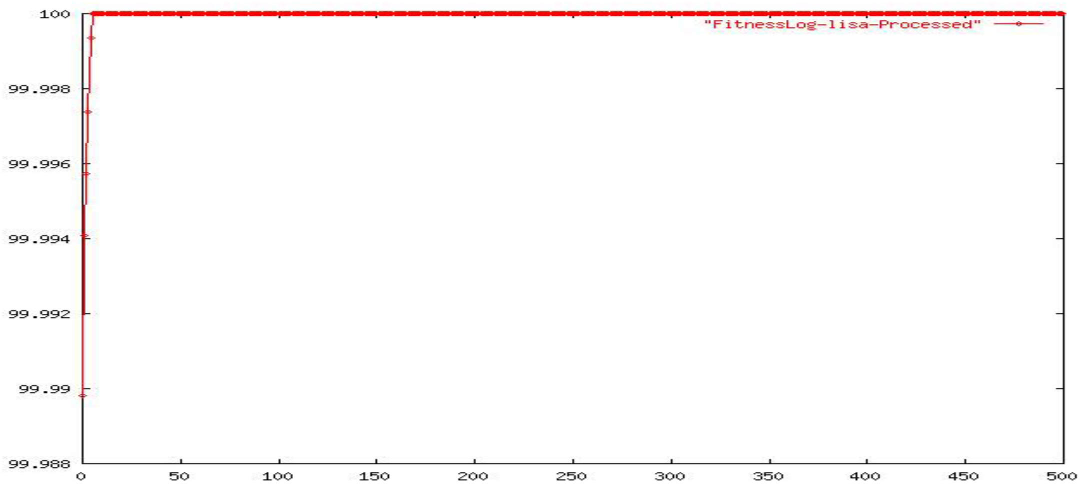

Fig 3.24 This graph shows elite fitness values, plus and minus one standard deviation, averaged over 25 runs.

For cases like SDIII, when we amplify the fitness landscape, depicted in Figure 3.25, it is reassuring to find that the underlying trend does seem to peak at around Qs=0.2 and confirms the result given by the ES.

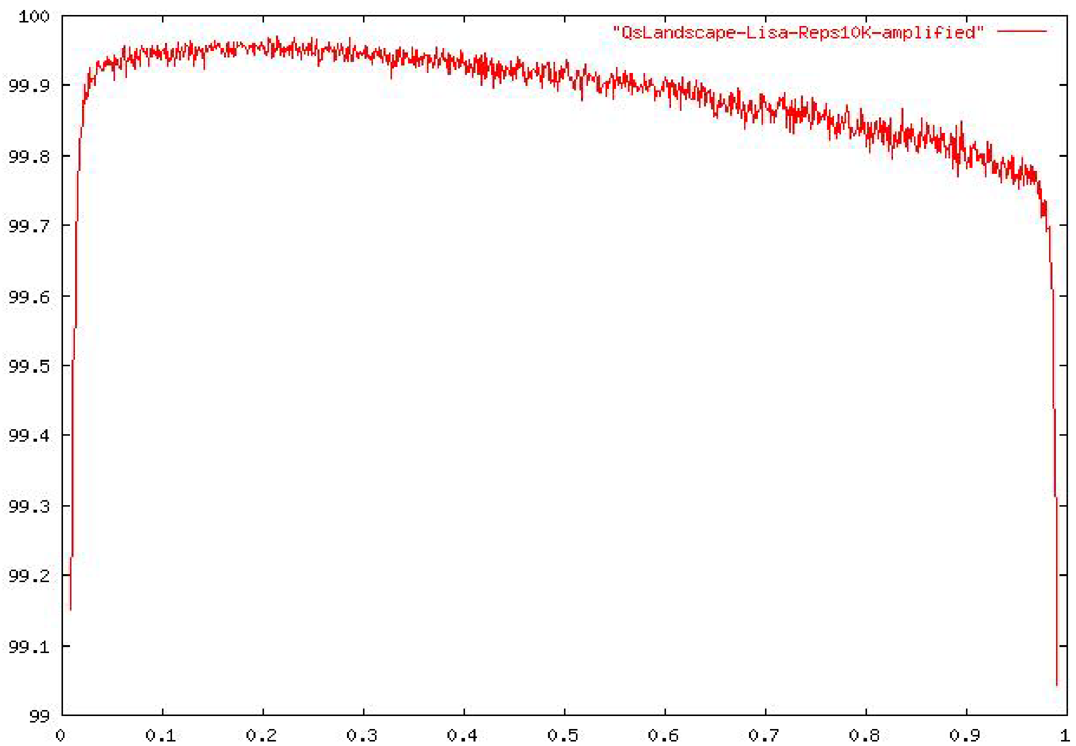

Fig 3.25 The figure shows the landscape for SDIII that has been rescaled and plotted for fitness values above 99%

3.3.6 M1

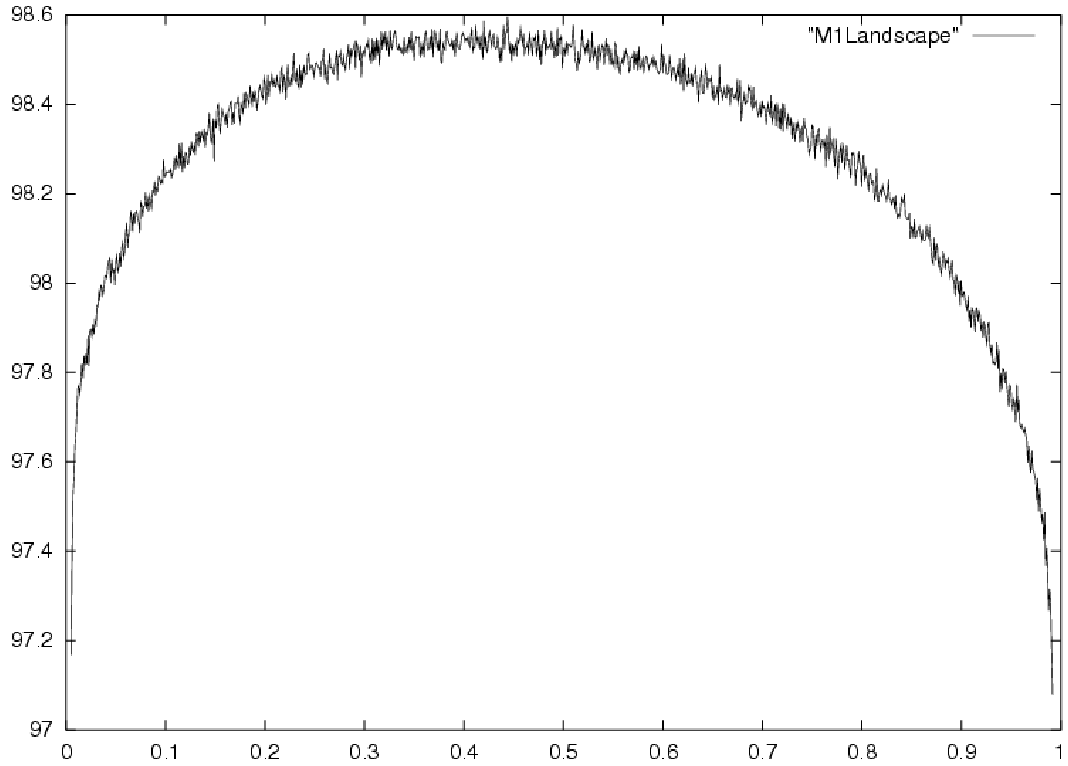

Fig 3.26. The fitness landscape for the M1 problem. See text for discussion.

Cliff (2001a) reports that his GA found a value of Qs=0.0001 to be optimum for market M1. It is obvious from the M1 fitness landscape shown in Figure 3.26 that this disagrees with the value of Qs that is optimum in a ZI-C populated market. We observe from Figure 3.28 that the elite Qs value averaged over 10 runs is around 0.45. This seems to agree with the landscape shown in Figure 3.26.

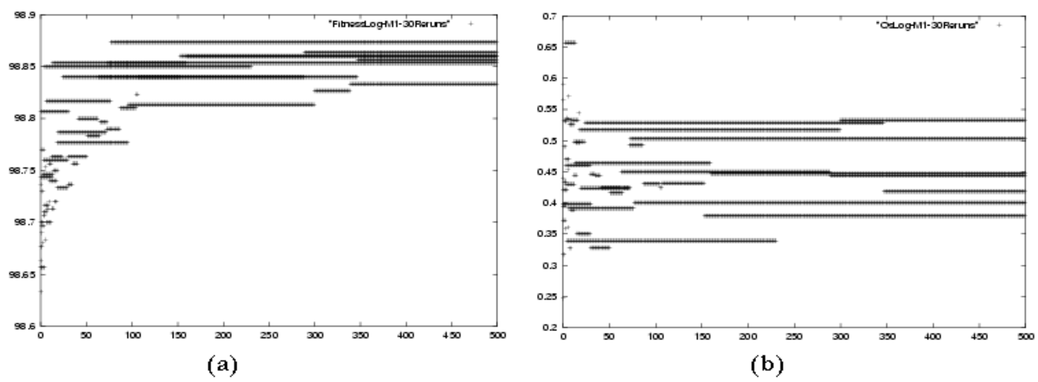

Fig 3.27 (a) Elite Fitness and (b) Elite Qs value logs for all 10 runs of the ES

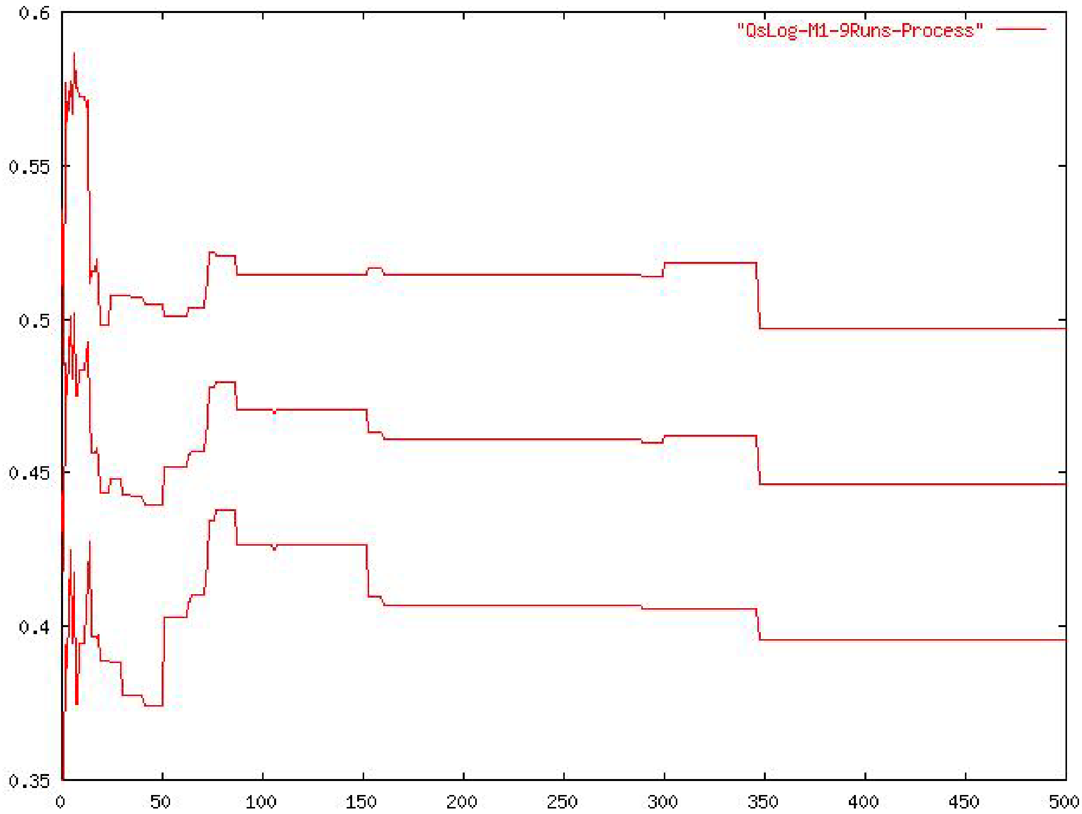

Fig 3.28 Elite Qs values averaged over 10 runs of the ES. The plots for plus and minus one standard deviation are also shown

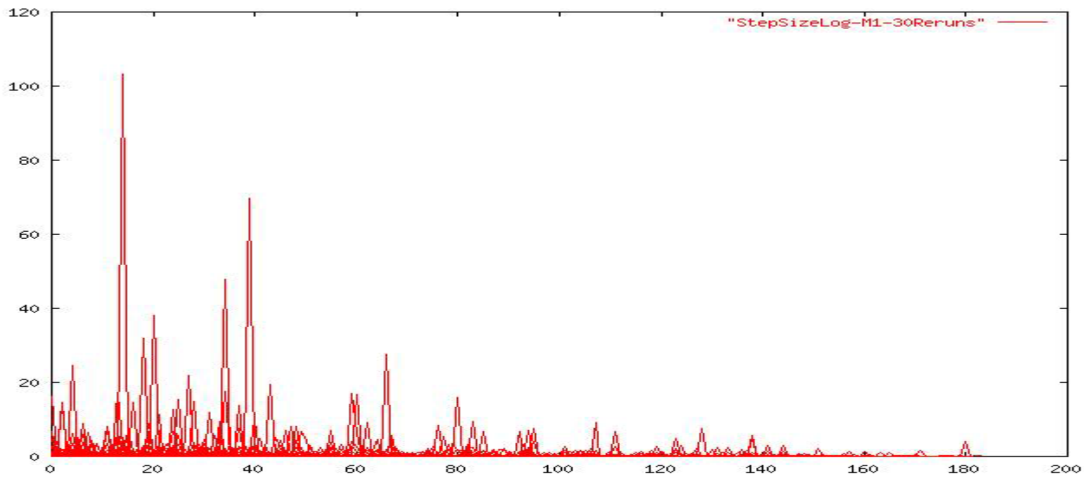

Fig 3.29 Step Size Log for all 10 runs of the ES on the M1 problem

We observe from Figure 3.29 that there is a steady decrease in step size as evolution proceeds. Figure 3.30 shows the elite fitness values averaged over the 10 runs of the experiment.

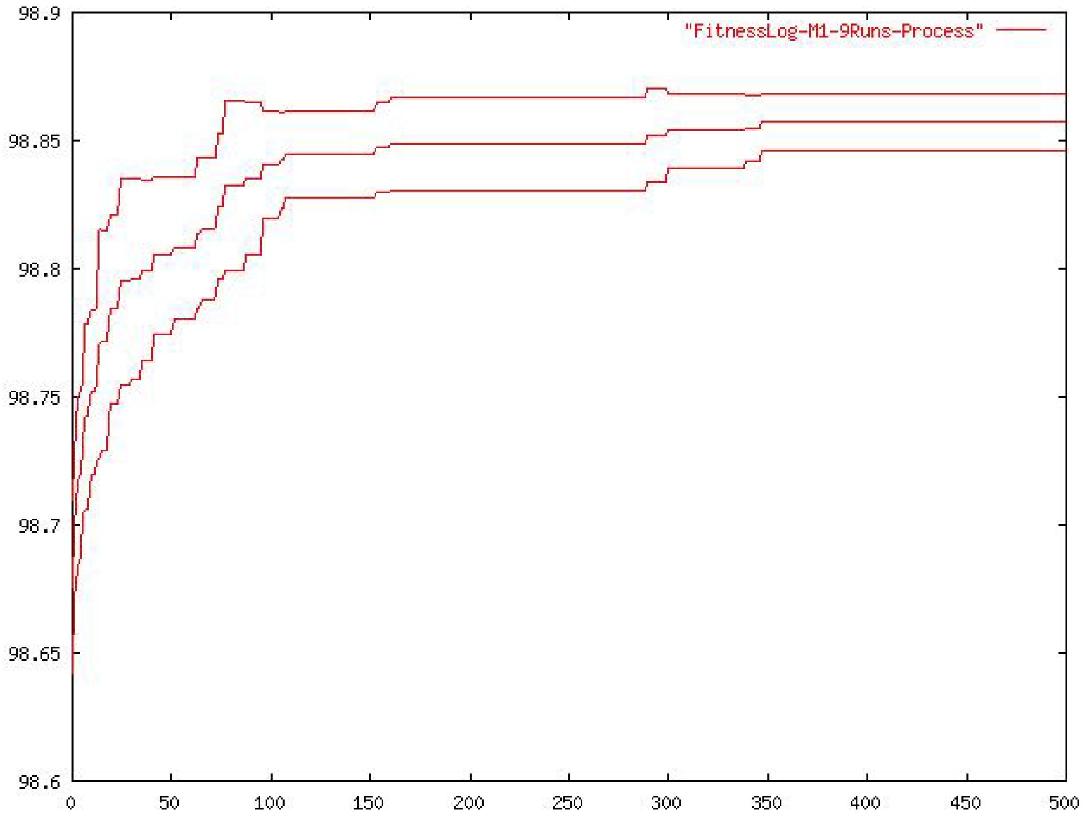

Fig 3.30 Elite Fitness log averaged over 10 runs of the ES on problem M1 with plus and minus one standard deviation

3.3.7 M2

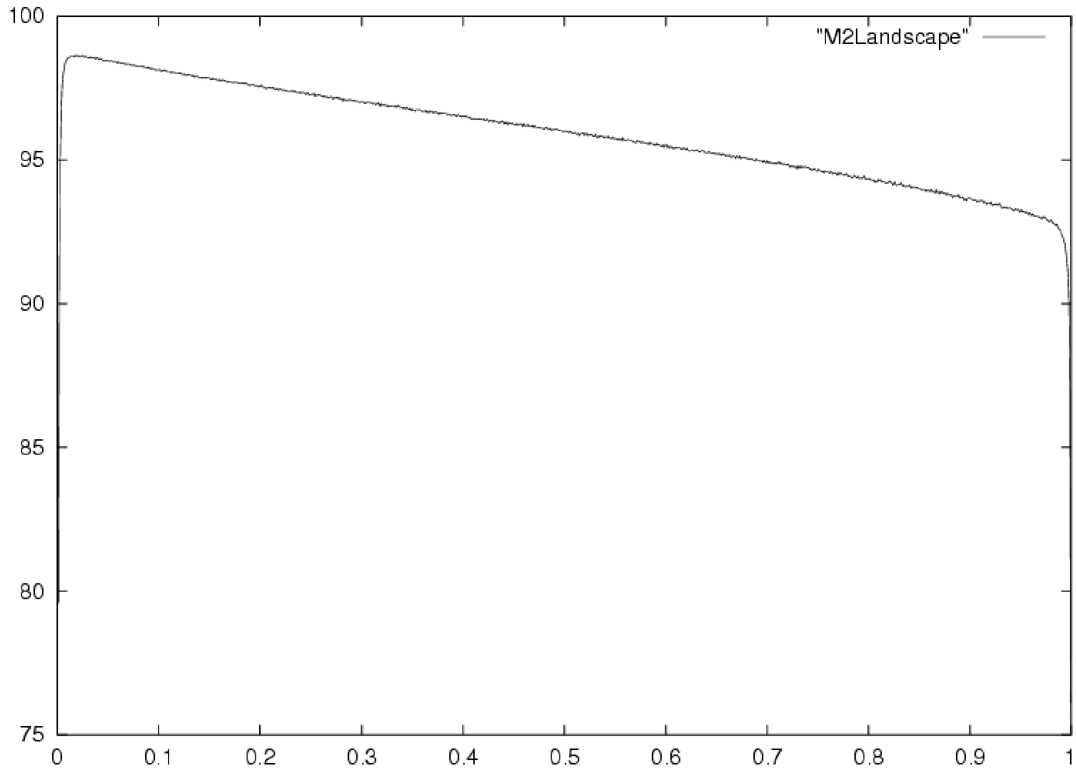

Fig 3.31. The fitness landscape for the M2 problem. See text for discussion.

Figure 3.31 shows an almost linear monotonic increase in fitness as Qs varies from 1→ 0. But we notice that there are sharp 'stilts' at either end which seem to break the trend that is followed at values except those close to 0 and 1. This is because of the auctioning method that we have followed. Due to the presence of a queue in the persistent double shout version of the CDA, a trade occurs only when two compatible quotes cross. Since at Qs=0, the sellers cannot quote offers, trades will occur only in the case where the initialising offer made by the seller at the start of trading meets a bid by the buyer. We think that this edge effect is an artefact caused by the combination of both the auctioning mechanism we have used and the ZI-C trading algorithm (which initialises the starting bids to random values).

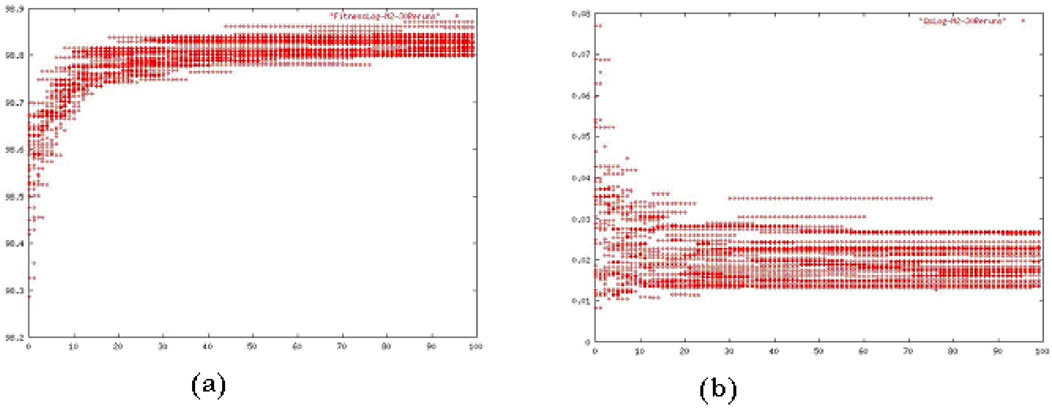

Fig (a) Elite Fitness and (b) Elite Qs value logs for all 10 runs of the ES on the M2 problem

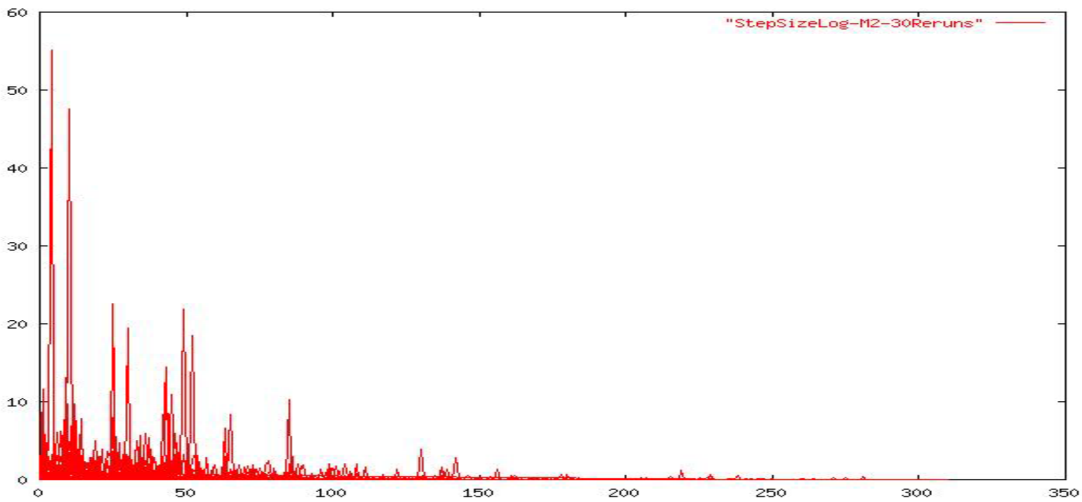

Fig 3.33 Step Size Log for all 10 runs of the ES on problem M2

We note from Figures 3.32 a and b, that the ES shows strong convergence in market M2. This is further confirmed by the steady decay of the step size as evolution proceeds.

Figure 3.34 shows that the evolved optimum Qs value is 0.02. Cliff (2001a) reports an optimum Qs value of 0.07. We think that save for the edge effect, the optimum value of Qs in this market is actually very close to 0. This indicates that in a market like M2, an English Auction like, one-sided mechanism would be optimum. This demonstrates that there are markets in which one-sided mechanisms can emerge to be optimum and that our ES can successfully indicate the same.

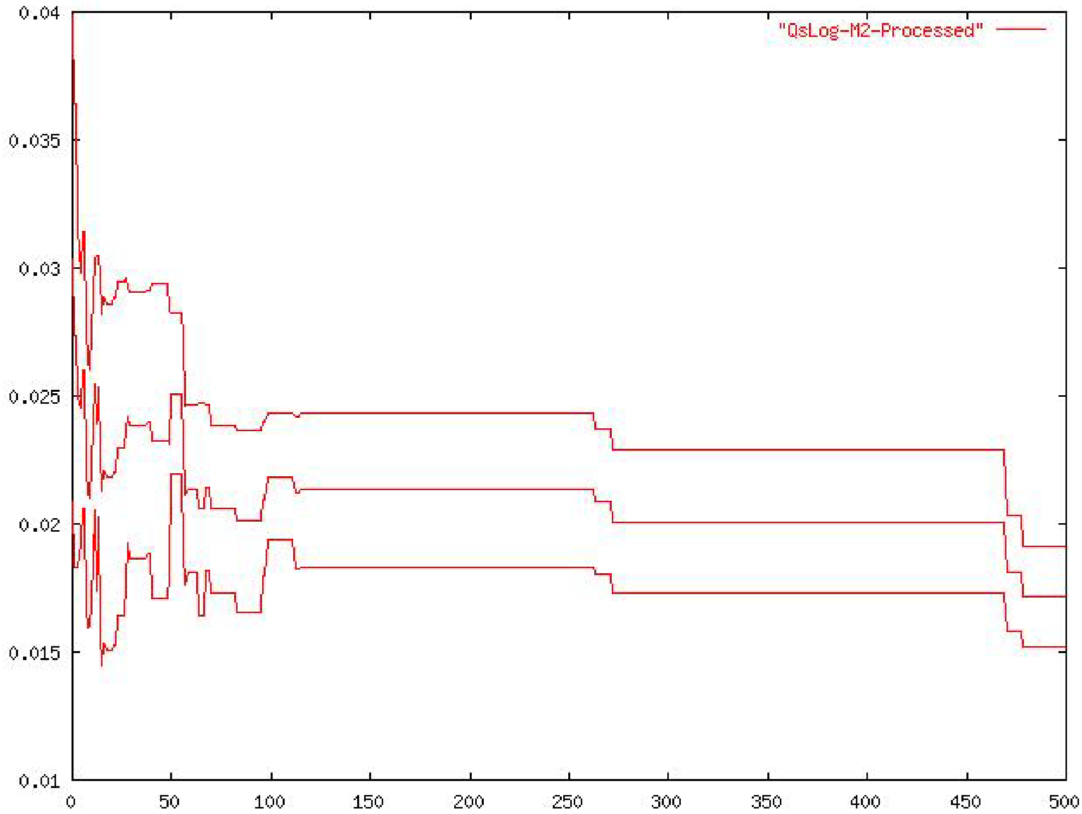

Fig 3.34 Elite Qs values averaged over 10 runs of the ES for problem M2 plus and minus one standard deviation

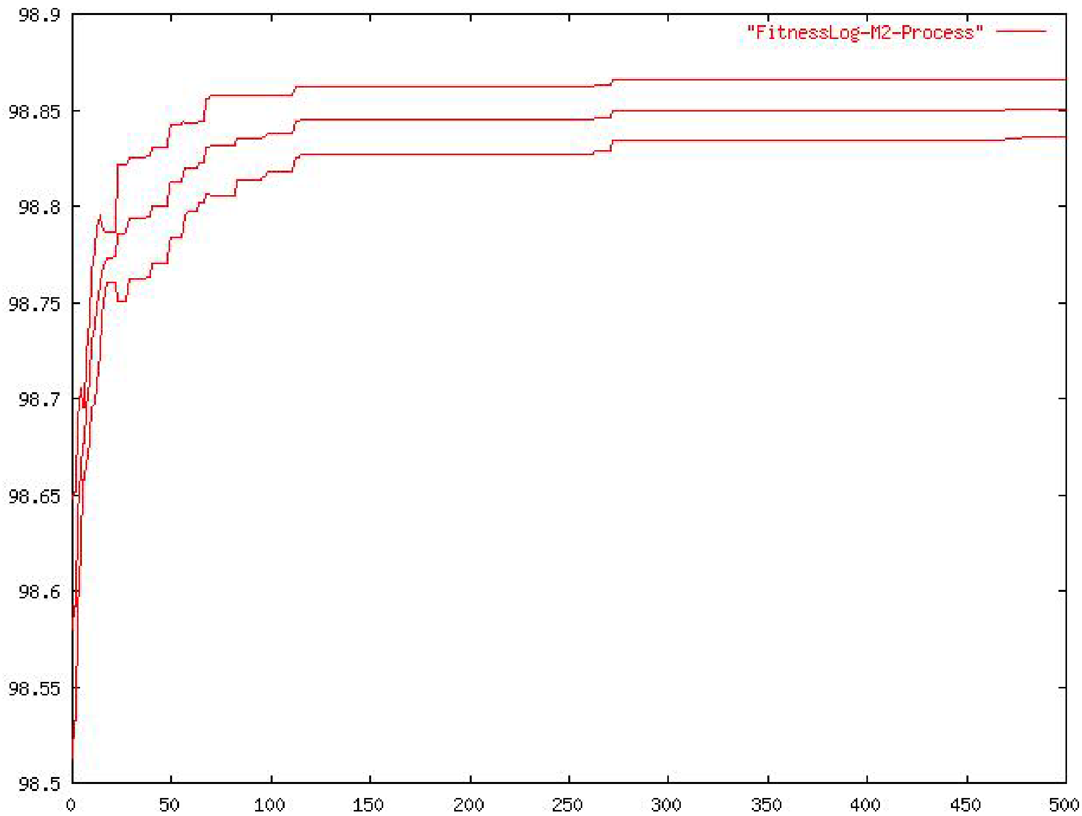

Fig 3.35 Elite fitness values averaged over 10 runs of the ES for problem M2 plus and minus one standard deviation

3.3.8 M3

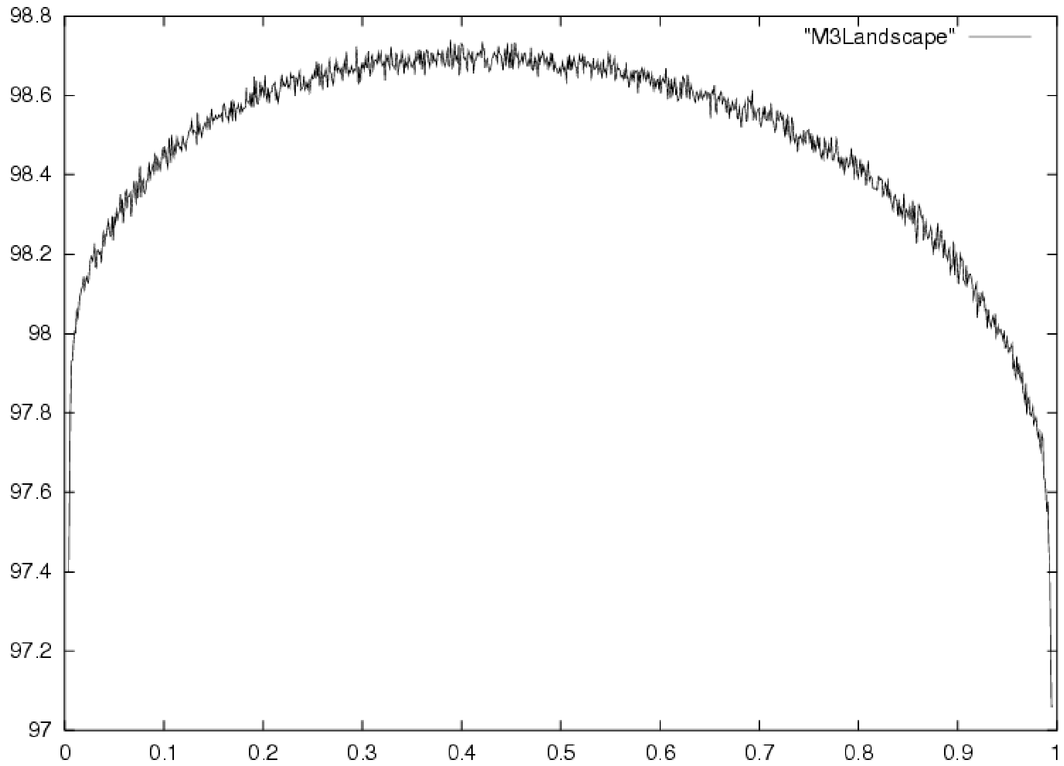

Fig 3.36. The fitness landscape for the M3 problem. See text for discussion.

Cliff (2001a) reports an optimum of Qs=0.16 for market M3. Once again, this disagrees both with the landscape shown in Figure 3.36 as well as the average elite Qs value shown in Figure 3.38. Figure 3.38 shows that the optimum Qs value lies close to Qs=0.4.

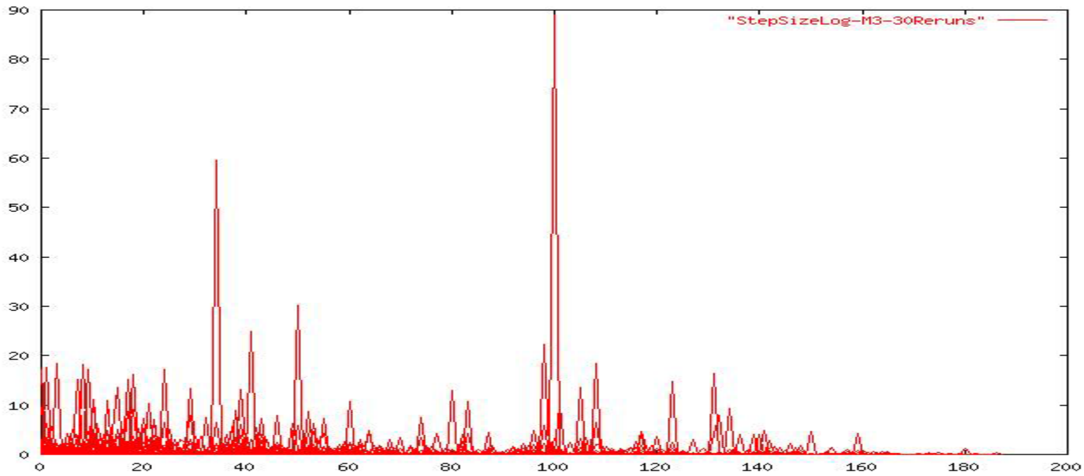

Fig 3.37 Step Size log for all 10 runs of the ES on problem M3

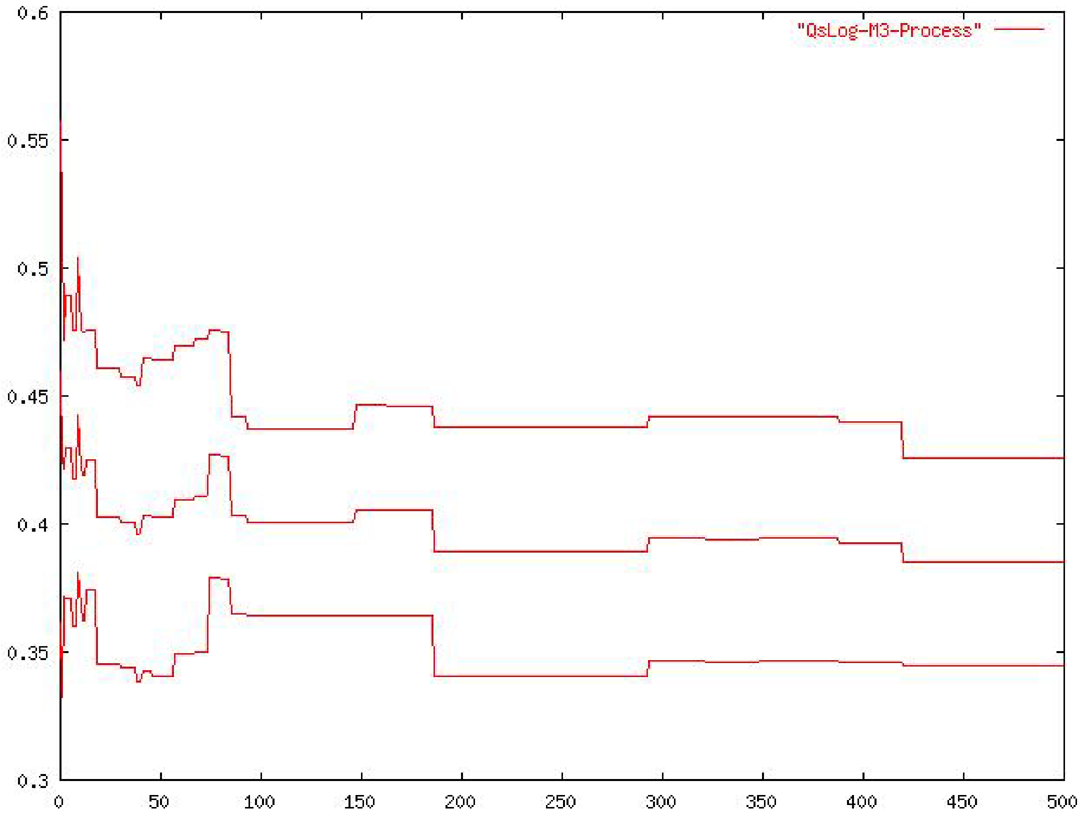

Fig 3.38 Elite Qs values averaged over 10 runs of the ES for problem M3 plus and minus one standard deviation

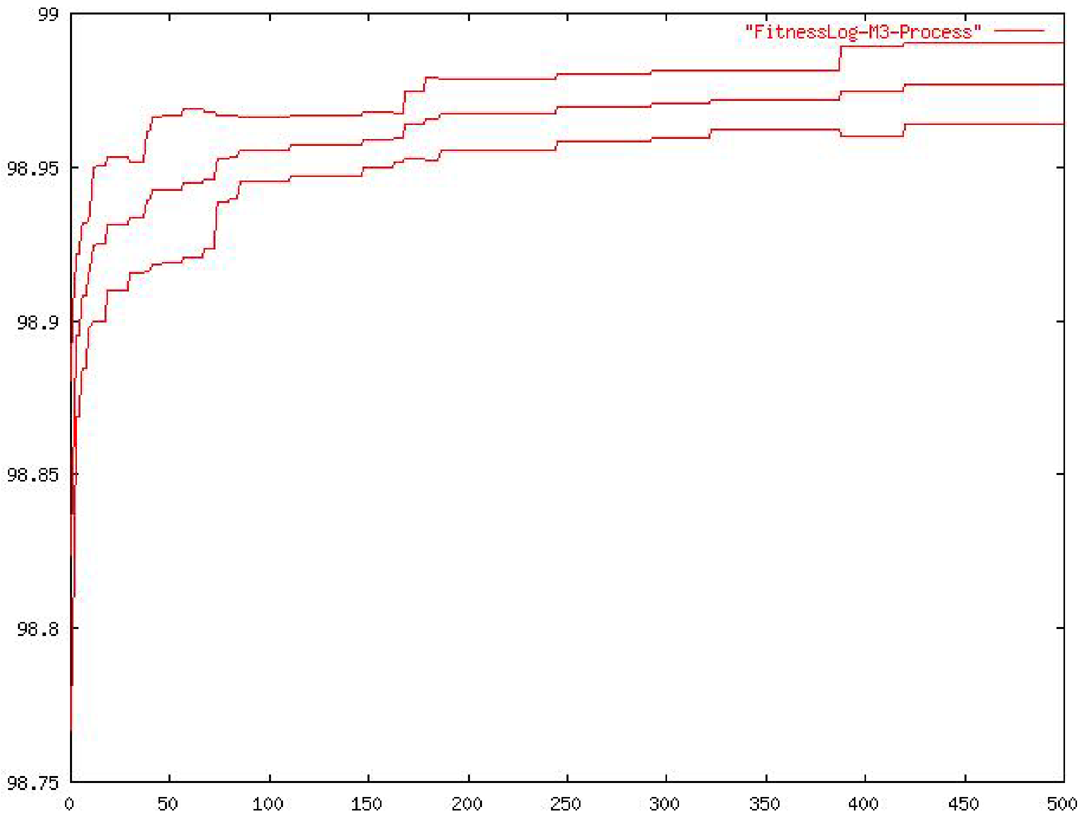

Fig 3.39 Elite Fitness values averaged over 10 runs of the ES for problem M3 plus and minus one standard deviation

3.4 Confirming ES Results

Given the extremely noisy nature of the fitness evaluation function and the stochasticity that is associated with any evolutionary system how can we check the validity of the results that are reported by our ES?

One was to do this is to generate exhaustive landscapes, with fitness values evaluated many time over ($\theta(10^6)$). However, as indicated before exhaustive landscapes can be prohibitively expensive to generate. A landscape plotted for 1,000 values of Qs with fitness values averaged over 100,000 repeat evaluations can take up to 10 days to generate. However, once the ES indicates which regions are of special interest we can exhaustively 'landscape' points of interest and observe whether the optimum Qs value reported by the ES truly dominates points in its neighbourhood.

We perform a series of such tests for problems M1 to M3 and show that the average fitness value for the points indicated by the ES to be optimum consistently dominates neighbouring points. We evaluate fitness for points of interests by averaging over a large number of runs ($\theta(10^6)$) and repeat the experiment multiple number of times to ensure greater reliability.

3.4.1 M1

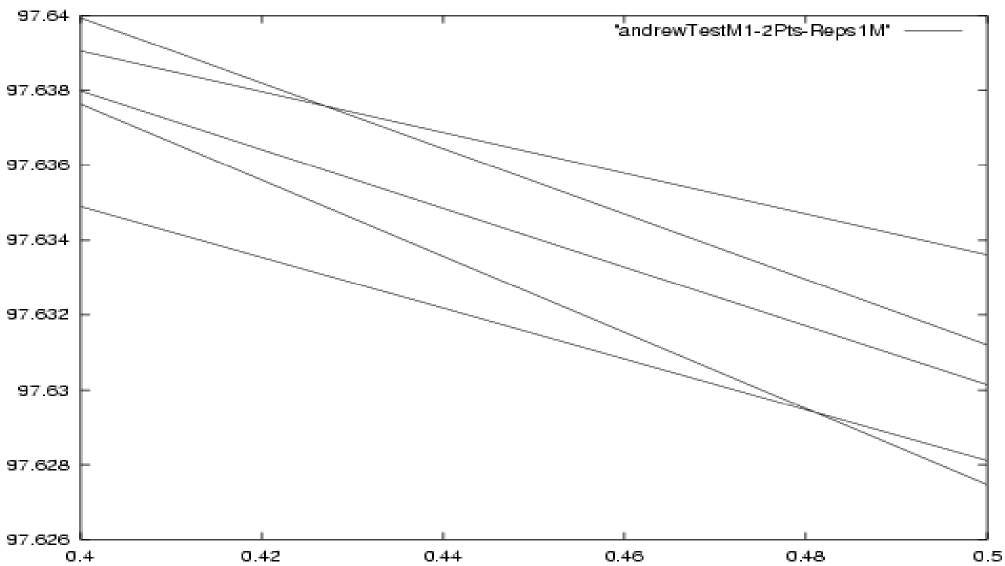

Fig 3.41 Fitness Values for Qs=0.4 and Qs=0.5 are averaged over 1 million repeats and plotted to show that on an average Qs=0.4 dominates Qs=0.5 w.r.t. fitness. The experiment is repeated 5 times.

We observe from Figure 3.40 that Qs=0.4 dominates Qs=0.5. Further we note from Figure 3.41 that Qs=0.4 dominates Qs=0.3 as well. Note that the difference in average fitness is ~0.01%.[13]

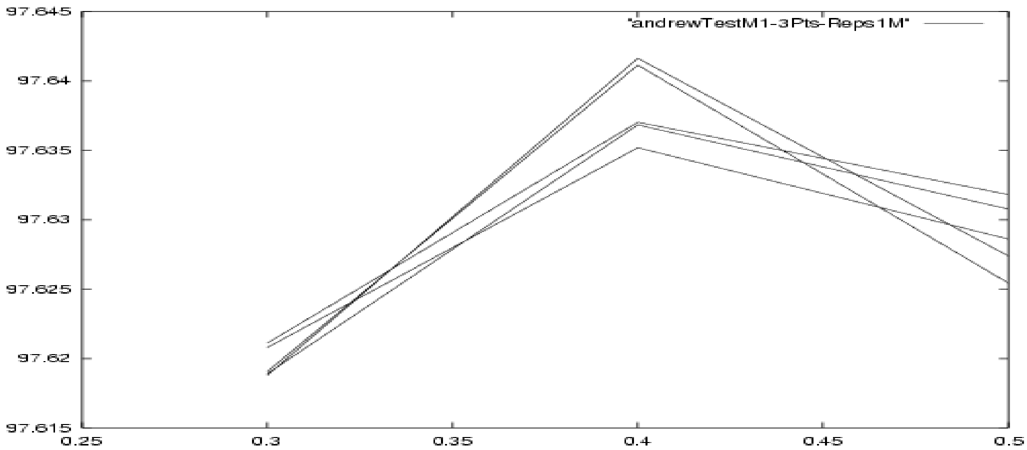

Figure 3.41 Fitness Values for Qs=0.3, 0.4 0.5 are averaged over 1 million repeats and plotted to show that on an average Qs=0.4 dominates 0.3 and 0.5 w.r.t. fitness. The experiment is repeated 5 times.

[13] This difference is significant if the absolute actual value of trades in the market is large. On a typical trading day the New York Stock Exchange (NYSE) witnesses more than 20 million trades with financial transactions valued at over 33 billion dollars. (Source: NYSE official statistics for 20th of August 2002 from NYSE's official website http://www.nyse.com)

3.4.2 M2

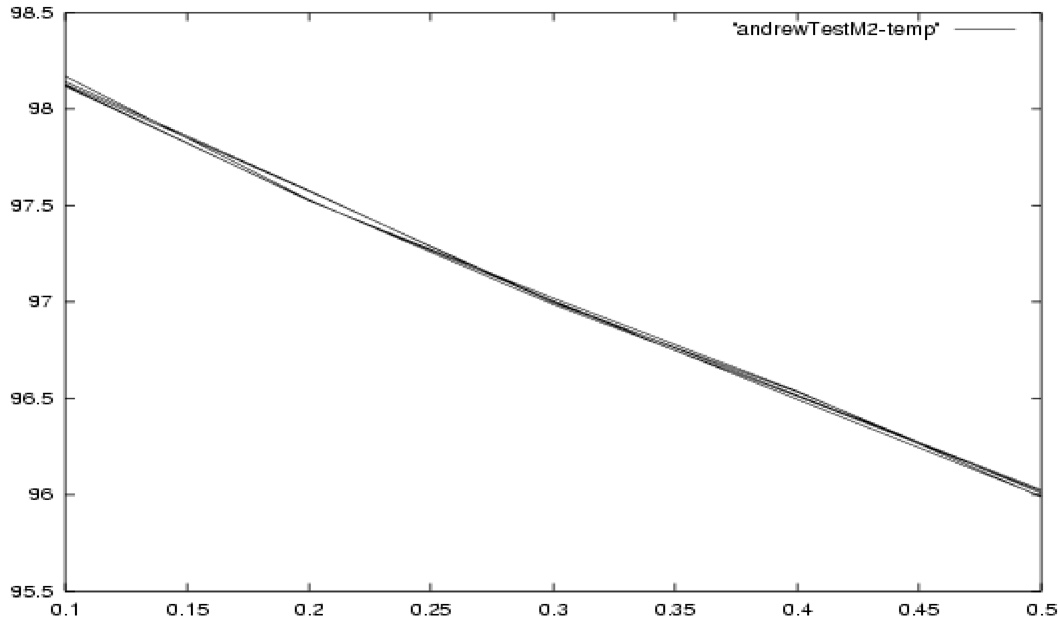

Fig 3.42 Fitness Values for Qs=0.1 and 0.5 are averaged over 1 million repeats and plotted. The experiment is repeated 5 times.

A value of Qs→0 (but not actually 0) is indicated to be optimum by the ES. Figure 3.42 helps confirm this broad trend and more importantly helps shows that the difference in the average efficiency obtained by organising the market as a Qs=0.1, as opposed to a conventional CDA (Qs=0.5) is larger than 2% on an average.

3.4.3 M3

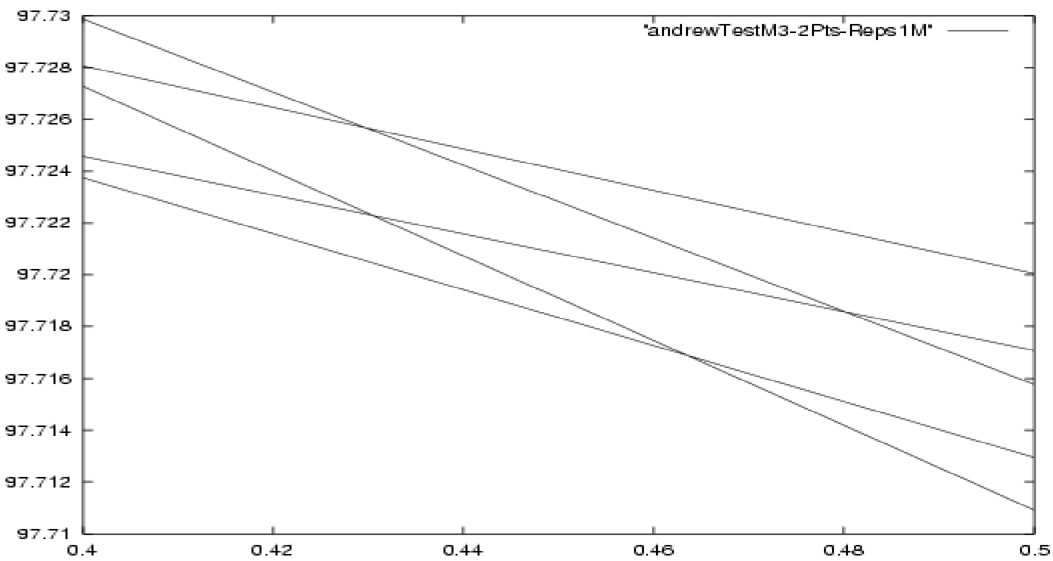

Fig 3.43 Fitness Values for Qs=0.4 and 0.5 are averaged over 1 million repeats and plotted to show that on an average Qs=0.4 dominates 0.3 w.r.t. fitness. The experiment is repeated 5 times.

We observe from Figure 3.43 that the value of Qs=0.4 indicated to be optimum by the ES dominates Qs=0.5 (conventional CDA) by around 0.01%. Moreover, Figure 3.44 shows the coarse (in terms of points plotted) trend and confirms that Qs=0.4 dominates other points in its neighbourhood as well.

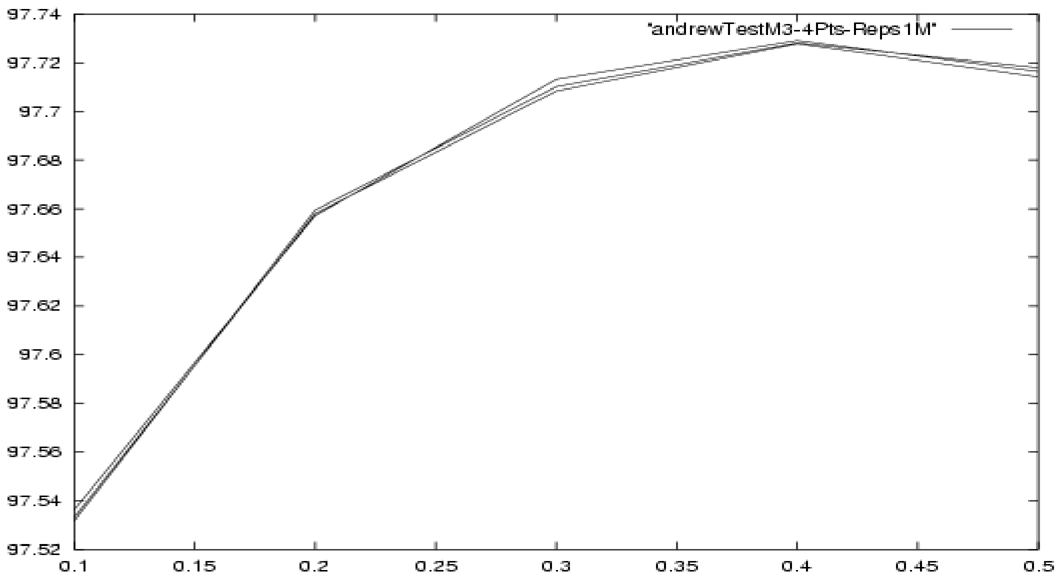

Figure 3.44 Fitness Values for Qs=0.1,0.2,0.3,0.4 and 0.5 are averaged over 1 million repeats and plotted to show that on an average Qs=0.4 all other pts w.r.t. fitness. The experiment is repeated 5 times.

3.5 Summary

We first described our experimental set-up in Section 3.1. We explained why we used ZI-C traders and highlighted the rationale behind some of our design choices. We feel that the fact that hybrid variants of the CDA dominate in spite of the changes we have made to the actual mechanism implemented is strong support for the claim that in general hybrid variants can lead to larger gains.

We described the design of our Evolutionary Strategy in detail in Section 3.2 and gave results from a set of 6 experiments in Section 3.3. Our results have show that hybrid variants can be superior to conventional variants in certain markets. However, there still are some markets in which conventional variants of Qs=0, 1 or 0.5 might dominate. We feel that this supports our notion of these conventional mechanisms simply being points on a continuum and shows that our ES has no a priori bias towards finding only hybrid variants.

Chapter 4

Gaining Insight in to Empirical Results

The results given in the previous chapter offer strong empirical evidence for the claim that asymmetric Double Auctions (DA) offer larger gains than are possible with conventional market mechanism types. In this chapter we report results from some experiments that we have conducted to gain insight into the relationship between the optimum Qs value for a given market and the nature of the underlying supply and demand curves.

The methods we have described in the previous chapter can help us determine the optimum value of Qs for a given market, which can then be used to conduct an auction which offers a high likelihood of making larger gains than previously known auction types like the English, Dutch and Continuous Double Auction.

However, our technique requires: -

1. A perfect or good knowledge about the nature of the underlying Supply and Demand curves for the market.
2. Significantly large periods of time (of the order of hours) to conduct an evolutionary search through the space of possible market types

An argument against the applicability of such a technique is that the underlying supply and demand data for any given market is unknown and moreover, changes dynamically with time. This is a fair criticism and is perfectly applicable to volatile markets in which the underlying supply and demand curves change rapidly. However, there exist markets which are relatively stable and whose nature changes only slowly with time. In these markets it is possible to determine values of Qs that can be used for protracted periods of time.

However, it would be best if we could determine the optimum Qs value in spite of imperfect supply and demand data and in a relatively short period of time. Evolutionary techniques are probably ill suited for dynamically determining the optimum Qs value in a 'live' market. So in order to be able to determine the optimum Qs in a shorter period of time we need to gain more insight in to the relationship between the optimum Qs value and the underlying supply and demand curves.

4.1 Previous Work

Gode and Sunder (1992) describe work in which they "attempt to determine the lower-bounds for expected surplus extraction of an idealised double auction populated with budget constrained ZI-C traders". They state: -

> "The double auction appears to be too complex to yield in a clear game-theoretic solution; its properties are easier to analyse with traders…"

This is the only work we know of in which there has been attempt to gain insight into the relationship between the underlying supply and demand curves for a given market and the expected outcomes using an agent-based approach. Gode and Sunder examine two variants of the Double Auction mechanism and present both analytical and empirical arguments to prove that the lowest expected efficiency for any given market organised as a Double Auction is 75%. But the model they have used to describe the underlying supply and demand curves is highly unrealistic and does not generalise well to actual laboratory or field institutions.

4.2 Experimental Set-up

We have decided to parameterise markets in terms of the slope of the supply and demand curves that describe them. We assume that the supply and demand curves are linear in nature. Moreover since the number of traders in our experiments is small (10-22) we must devise a way to define a discrete analogous of supply slope and demand slope.

Our markets are defined by a set of five parameters: -

1. **minPrice:** minimum price in the market
2. **maxPrice:** maximum price in the market

3. **mS:** the slope of the supply curve
4. **mB:** the slope of demand curve
5. **nB:** number of buyers in the market
6. **nS:** number of sellers in the market

The Supply Curve i.e. the limit prices for sellers are generated as follows: -

- The first seller s_0 is assigned price minPrice.
- Every subsequent trader s_i is allocated a price (minPrice+i*mS) until (minPrice+i*mS)>maxPrice.
    - If (minPrice+i*mS)>maxPrice the limit price for seller s_i is fixed at maxPrice.

Similarly for the demand curve: -

- The first buyer b_0 is assigned price maxPrice.
- Every subsequent trader b_i is allocated a price (maxPrice-i*mS) until (maxPrice-i*mS)<minPrice.
- If (maxPrice-i*mS)<minPrice the limit price for seller s_i is fixed at minPrice.

This procedure defines how we interpret the concept of slope for a discrete number of sellers. Figure 4.1 shows the supply and demand schedule for the following parameter set: minPrice=0, maxPrice=20, mS=10, mB=2,nB=5 and nS=5.

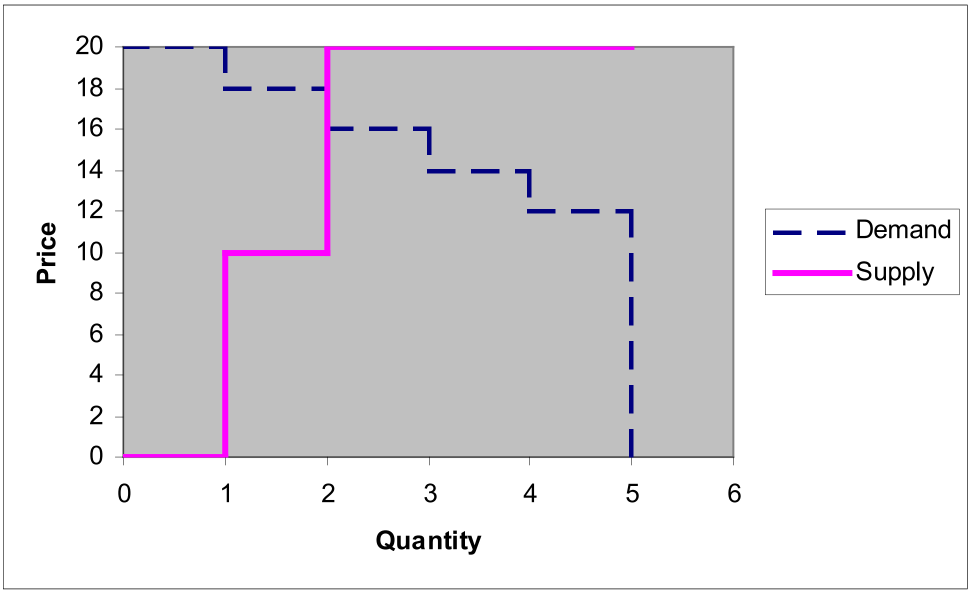

Buyers	20	18	16	14	12
Sellers	0	10	20	20	20

Fig 4.1 The supply and demand schedule generated for minPrice=0, maxPrice=20, mS=10, mB=2,nB=5 and nS=5

4.2.1 Algorithm

We use the ES discussed in last chapter to optimise Qs values over a large number of markets which can be interpreted as points on a (mB x mS) grid. We step though this matrix by generating a test market for each value (mB, mS) and run our ES on it to find the optimum value for this point on the grid.

If we vary the slopes through 21 points ([0,20]) each of mB and mS, we have 441 markets in which to need to find the optimum Qs value. If we set MAX_GENS to 500 and run the ES to completion for each of these markets this experiment would take more than 3 months to perform on a single machine. To gain computational leverage in this large landscaping process we have made some modifications to our ES. These modifications are based on the observations that very often the ES can stagnate much earlier than the maximum number of generations. The basic idea to is to monitor the behaviour of the ES as it optimises any market (mB,mS) and terminate its execution if the ES stagnates. We define a parameter MAX_STAGNANT and terminate if the ES has been stagnant for the last MAX_STAGNANT generations. Note that this parameter can have a significant effect on both the reliability of the test results as well as the computational leverage gained. Lower values result in faster execution but higher values result in a cleaner underlying landscape.

There are two criterions for judging if the ES has stagnated.

1. Elite Qs value has stagnated

First, we can terminate the execution of the ES if the elite Qs value has stagnated as this indicates a likely convergence to the optimum Qs value. We have included a bias in our ES that leads to an offspring replacing a member of the parent population if the fitness of both is equal. So if the elite Qs value has remained constant – this is probably on account of domination of this value over others rather than a random walk by the ES, which can happen in case of flat (or nearly flat) landscapes. Hence, in this case we terminate the ES and report the elite Qs value.

2. Elite Fitness value has stagnated

The case in which the fitness value stagnates, but the elite Qs value keeps changing is trickier. The reader should satisfy herself that there are landscapes in which the efficiency will always be 100%. Consider for example, the market for minPrice=0, maxPrice=20, mS=2, mB=2,nB=5 and nS=5.

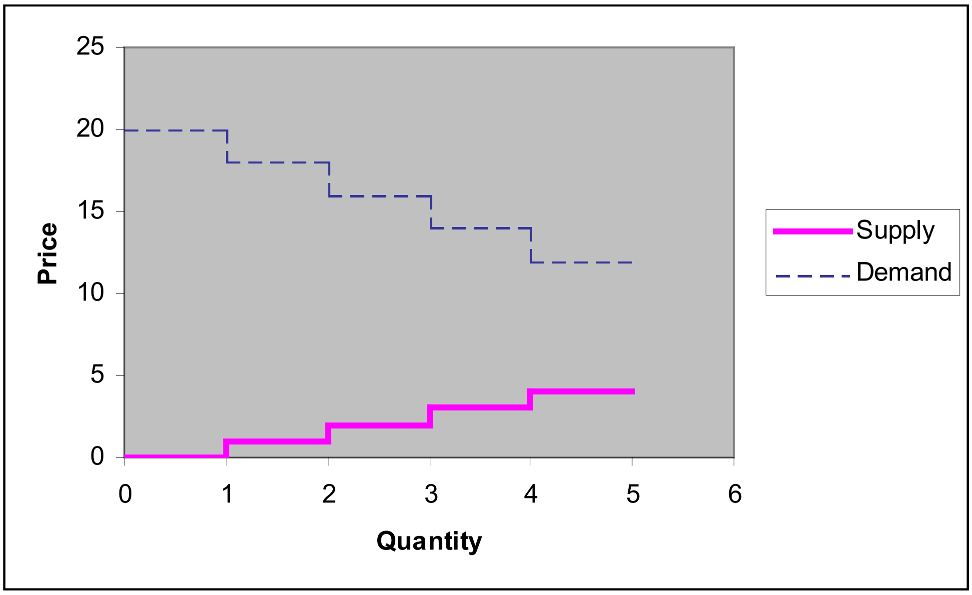

Fig 4.2 The Supply and Demand curves in this market do not intersect. As a result all traders in the market are intra-marginal and 100% efficiency is guaranteed if all the traders execute trades by the end of the trading period. Such markets will have very flat Qs landscapes with the baseline efficiency close to 100%.

The reader will observe that the supply and demand curves for this market do not intersect, meaning that all traders are intra-marginal leading to 100% efficiency for most runs regardless of the Qs value used. In such cases we have adopted the policy of reporting the average Qs value – which on an average can be expected to be ~0.5.

4.3 Results

We now report results for some experiments that we have performed.

Experiment 1

The mB and mS values for this experiment are varied over the interval [0,20] with an increment of 2. Thus we consider mB and mS = {0,2,4,6,8,10,12…20}. The number of buyers and sellers, nB=nS=5. Fitness values are averaged over 1,000 REPEAT_TRIALS. MAX-STAGNANT has a value of 20 for this experiment.

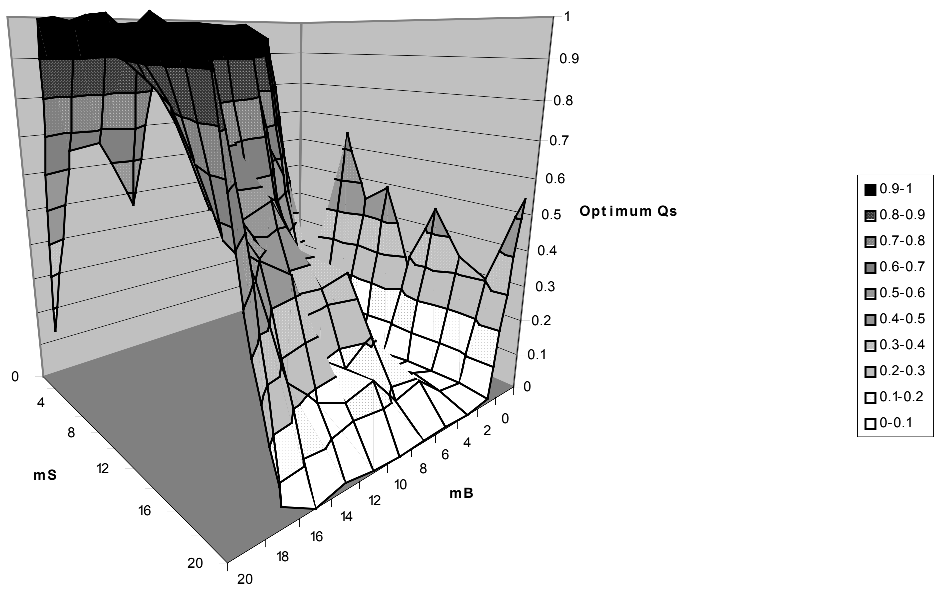

Fig 4.3 A 3-Dimensional landscape of optimum-Qs values over a large 2-dimensional space of test markets characterised by the slope of the Supply curve (mS) and Demand curve (mB). See text for discussion

The first thing that we note from Figure 4.3 is that the landscape is rotationally symmetrical across the diagonal. To gain a better idea of the actual distribution of Qs values we plot a 2-D contour plot (top view) for Figure 4.3 in Figure 4.4.

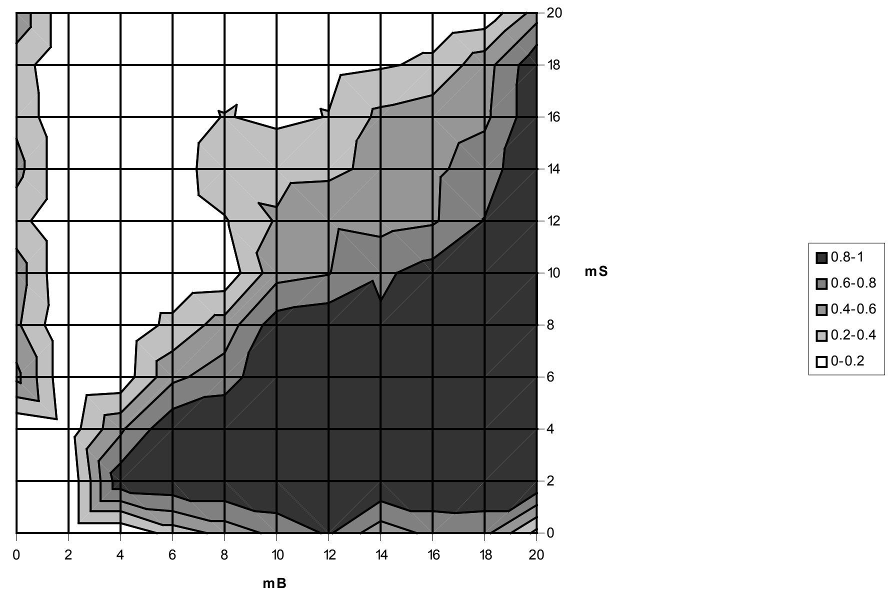

Fig 4.4 A contour plot of the 3-dimensional landscape shown in Figure 4.3. See text for discussion

We observe from the following regularities from Figure 4.4 that: -

1. Along the axis on which mB=20 almost all markets have an optimum Qs value of ~1.
2. Conversely when the slope of the supply slope mS=20 most markets have an optimum Qs value of ~0.
3. A significant number of markets have non-standard optimum Qs values.
4. We observe that along the y=x diagonal in this matrix, there is a consistent trend that which shows the domination of hybrid market types. Interestingly this non-standard (hybrid) Qs region (Qs∈[0.2,0.4], [0.4,0.6] and [0.6,0.8]) seems to broaden above the y=20-x diagonal in this matrix of markets.

We check the emergence of these trends by running another set of tests in which mB and mS values are varied over the interval [0,20] with an increment of 1. Thus we pick mB and mS = {0,1,2,3,4…20}. As in the previous test - the number of buyers and sellers, nB=nS=5. Also as in the previous test, fitness values are averaged over 1,000 REPEAT_TRIALS. However, MAX_STAGNANT has a value of 10 for this experiment (20 in last test). We make this change to gain more computational leverage, as the number of test markets in this case is 441 as opposed to 121 in the last experiment. The results from this experiment are shown below in Figures 4.5 and 4.6.

Experiment 2

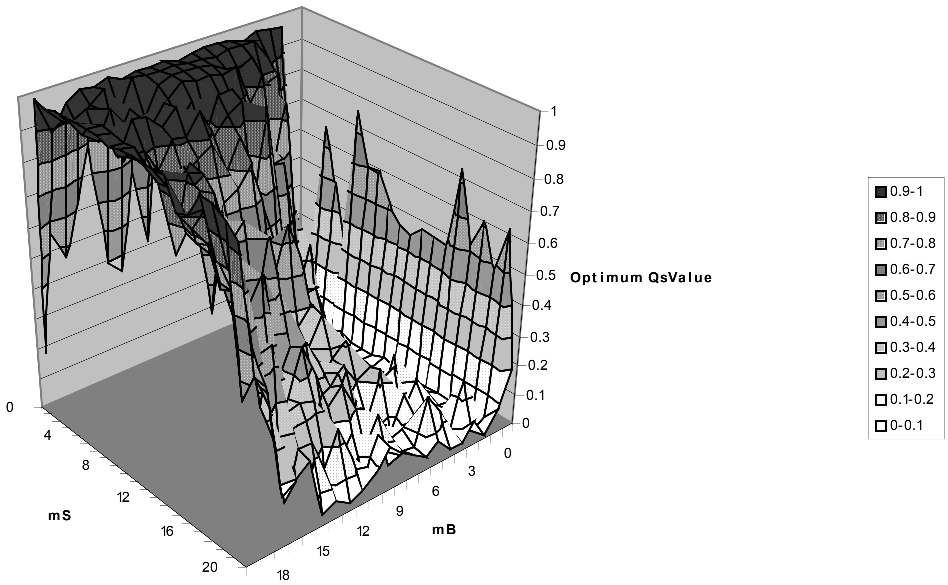

Fig 4.5 Three-Dimensional landscape of optimum-Qs values for Experiment 2. See text for discussion.

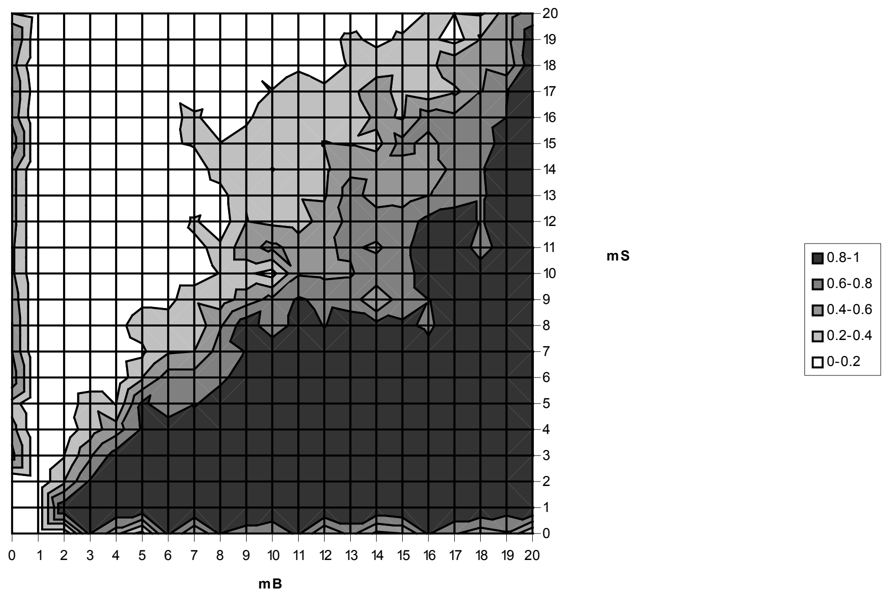

Fig 4.6 Contour plot of the 3-dimensional landscape shown in Figure 4.5. See text for discussion

We observe from Figures 4.5 and 4.6 that the trends outlined on Page 79 emerge again. We now try to provide qualitative explanations on what causes these trends to emerge.

Along the axis on which mB=20 almost all markets have an optimum Qs value of ~1

The Supply and Demand schedule for a market in which mB=20 is shown below. The value of mS is 2. But what is of essence in this market is the fact that due to the test-market generation procedure that we outlined on Page 75, there is only one intra-marginal buyer in this market.

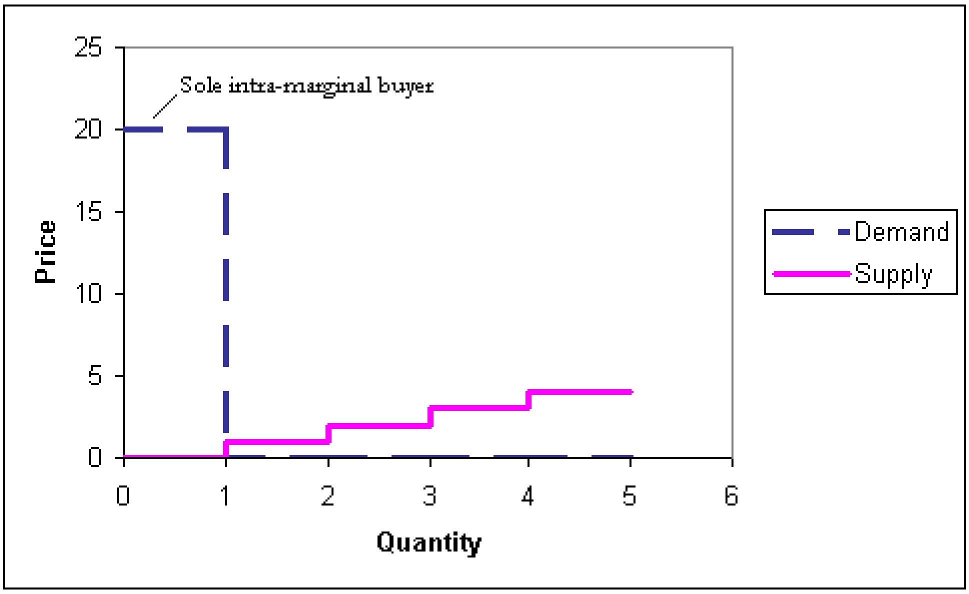

Fig 4.6 This figure shows the supply and demand schedule for a test market in which mB=20. Note that since there is only one intra-marginal buyer, ensuring 100% efficiency is a question of making sure that she trades with the seller with the lowest cost price.

For 100% efficiency to be achieved in this market we must make sure that the buyer with reserve (utility) price 20 trades with the seller whose cost price is 0, thus generating the largest possible total surplus of 20. If the intra-marginal buyer trades with any other seller, the surplus obtained will be reduced by an amount that is equal to the difference between the cost price c_i of seller i with whom the trade was executed, and the minimum cost price in the market c_{min}.

How does a Qs value of 1 help us in achieving this aim? A Qs value of 1 one means that only sellers shout offers in the market as the buyers observe silently[14]. This in turn means that sellers are being forced to compete amongst themselves and undercut each other continuously in an effort to steal the only trade that is available in the market. This

[14] This is similar to the conventional Dutch Auction.

competitive process leads to the elimination of the extra-marginal traders one-by-one as the current offer in the market drops below each of their reserve (cost) prices. This process continues till only the seller with the lowest reserve price is left in the market. This seller can now trade with the sole intra-marginal buyer generating 100% efficiency. So in this case a Qs value of 1 can be seen as an attempt by the mechanism to ensure that only the seller with the lowest cost price trades, thus leading to 100% efficiency.

Conversely when the slope of the supply slope mS=20 most markets have an optimum Qs value of ~0

The Supply and Demand schedule for a market in which mS=20 and mS=2 is shown below. As opposed to the market in the previous example there is only one intra-marginal seller in this market. The qualitative explanation given above for Figure 4.6 generalises to this example as well.

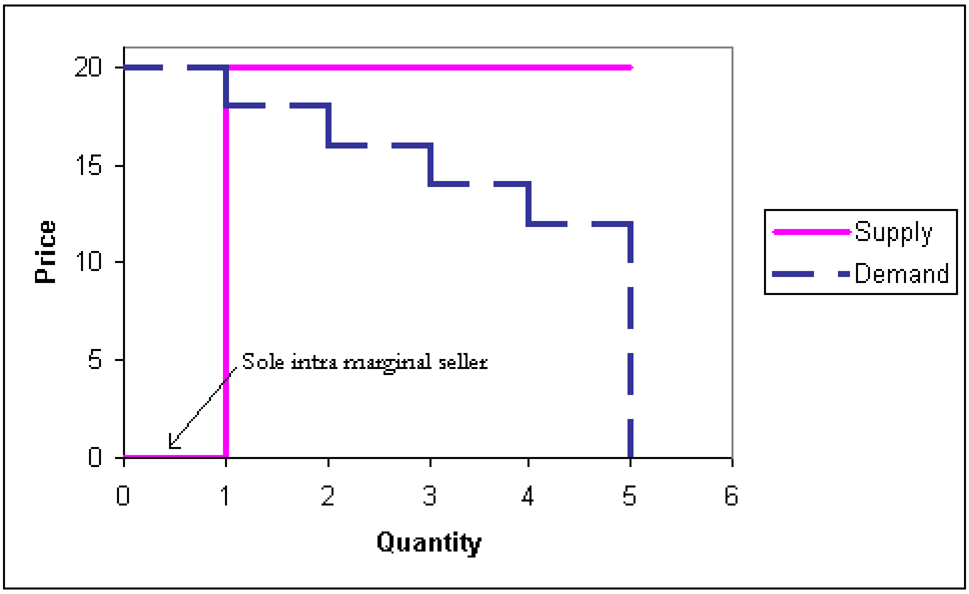

Fig 4.7 Supply and demand schedule for a test market in which mS=20. In this case there is only one intra-marginal seller, so achieving 100% efficiency is a question of making sure that she trades with the buyer with the highest reserve price (utility value).

Since there is only one intra-marginal seller but many buyers who are capable of executing trades in the market, achieving 100% efficiency in this case translates into ensuring that the intra-marginal seller trades with the buyer with the highest reserve price.

A Qs value of 0 means that all shouts in the market come from buyers. This leads to fierce competition among the buyers in the market which pushes the highest bid on the market up continuously till all the efficiency-reducing buyers have been pushed out of the market. So in this case a Qs value of 0 encourages the extraction of the maximum possible surplus from the market.

The two case discussed above lead to an interpretation of Qs as a factor that leads to skewing of the normal competition pressure that exists in a conventional CDA and makes the competition among either sellers or buyers fiercer. The value of Qs determines which side, the buyers or the sellers, is subjected to this increased competition. Values close to 0 increase the competition among buyers, values close to 1 among sellers. Furthermore the utility of this skewing emerges from the fact that the optimum Qs value for a given market can lead to an average gain in the surplus extracted.

This interpretation also leads to a better understanding of the effect the Qs value has on the auctioneering process and the realisation that there is no a priori reason for values close to 0,1 or 0.5 to be optimal for a given market. These are values represent markets in which there is competition only among sellers, only among buyers and equally among buyers and sellers but there may well be cases where some other 'mix' is optimal. These cases are represented by markets in which 'hybrid' values of Qs are found to be optimum.

We observe that along the y=x diagonal in this matrix, there is a consistent trend that which shows the domination of hybrid market types. Interestingly, this non-standard (hybrid) Qs region seems to broaden above the y=20-x diagonal in this matrix of markets

Let us calculate the Equilibrium Quantity Qo in terms of the slopes of the supply and demand curves mB and mS for the parameters we have used in our experiments (minPrice=0, maxPrice=20 and nB=nS=5).

Recollect that the equilibrium quantity Qo is given by the x co-ordinate of the point of intersection of the supply and demand curves. Hence,

$$mS.Qo + mB.Qo = 20$$

$$Qo = \frac{20}{mS + mB}$$

So for Qo=1

(mS+mB)>=20

For Qo=2

20>(mS+mB)>=10

For Qo=3

10>(mS+mB)>=6²/₃

For Qo=4

6²/₃>(mS+mB)>=5

And finally for Qo=5

5>(mS+mB)>=4

which is a market in which all traders are intra-marginal (for nB=nS=5)

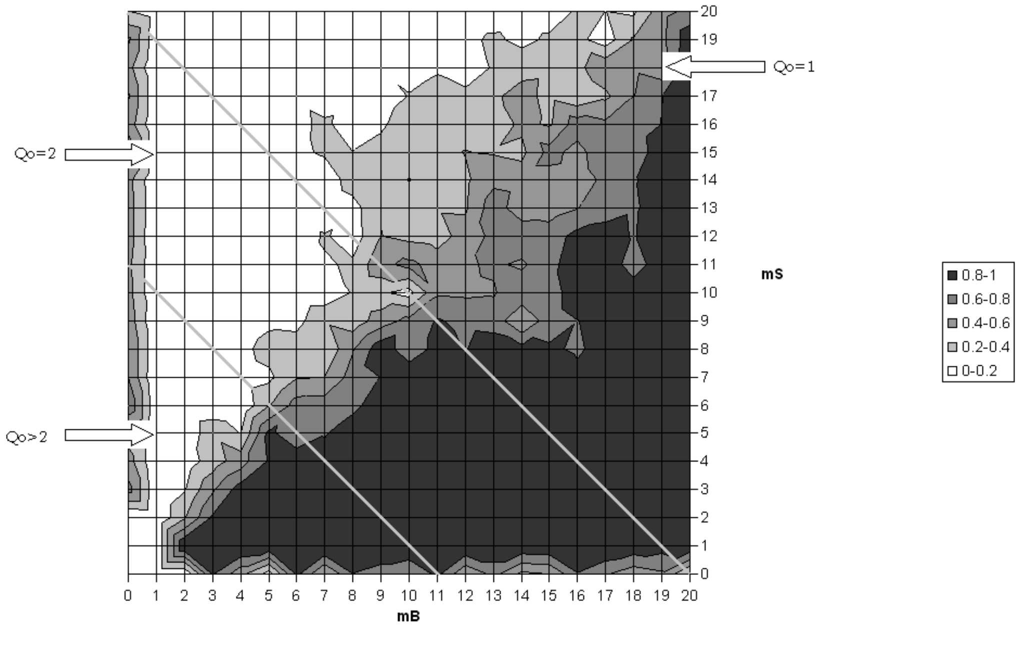

Fig 4.8 Contour plot for experiment 2 with the regions with Q_o values of 1,2 and greater than 2 marked. See text for discussion.

Examining the contour plot for experiment 2 we observe that hybrid values of Qs dominate more and more as the value of Qo decreases. This is shown by the broadening of the hybrid 'bulge' that can be seen along the y=x diagonal. This region is highlighted below in Figure 4.9.

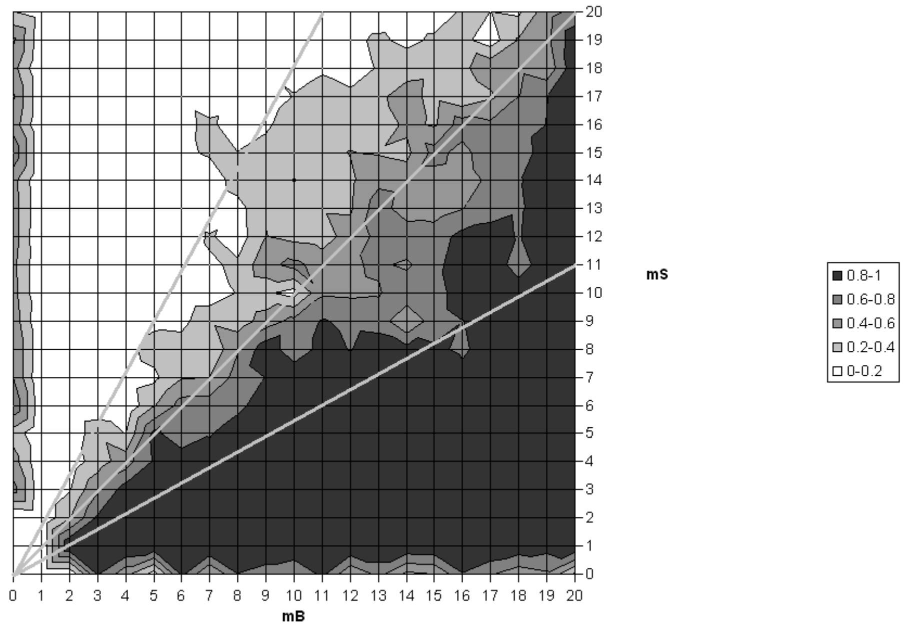

Fig 4.9 The contour plot for experiment 2 with the hybrid market dominated region marked. We observe that this region lies along the y=x diagonal and seems to get broader as the value of Q_o in the underlying test-markets increases.

As the value of Qo decreases, it becomes increasingly critical to get the only profitable trade available in the market 'right'. Since the efficiency achieved is directly related to the surplus generated in this *only* transaction, the effect of using Qs as a control parameter is heightened. If we look at Figure 4.8 again in this light we observe that the indeed, the 'hybrid' region does seem to broaden as Qo decreases.

We run another test over only this central 'hybrid' region of interest. Also we plot only half of this region since the market is perfectly symmetrical (along the y=x diagonal) with respect to the role of the buyers and the sellers. This can be observed from Figures 4.3, 4.4, 4.5 and 4.6. The contour plot for this data is shown in Figure 4.10. Another interesting observation is how this region seems to become narrower as mS increase further. Currently we do not understand why this happens.

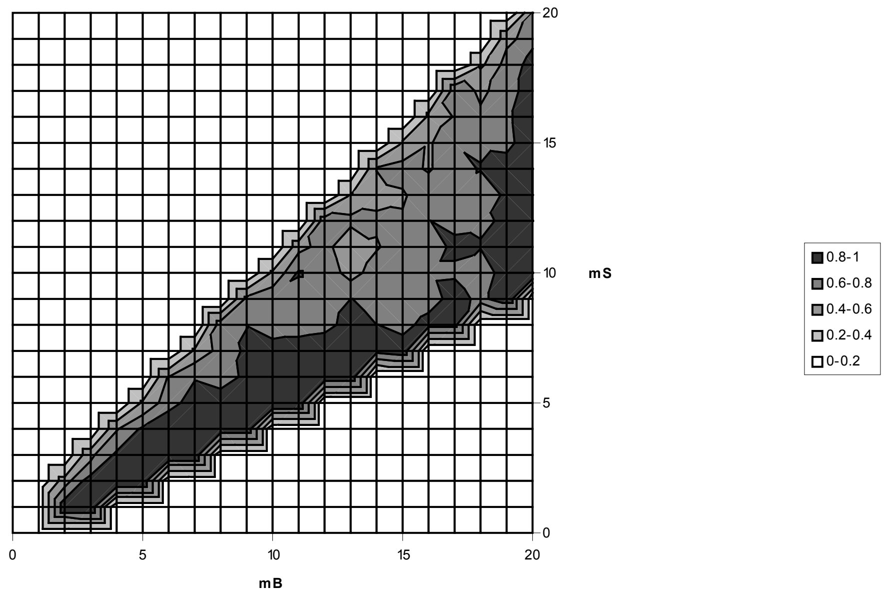

Fig 4.10 This figure plots the hybrid market dominated central region. Since the market is symmetrical w.r.t. the role of the buyers and sellers (i.e. along the y=x) diagonal only the lower half is plotted. Please note that the (0-0.2) region is actually not generated.

4.4 Summary

Having demonstrated in the last Chapter that hybrid markets can offer larger gains than more conventional variants, we try to explore the relationship between the supply and demand characteristics of a given market and the value of Qs that is found to be optimal for it. Such an understanding will be essential if we wish to apply this concept to real-time 'live' markets in which the nature of the underlying supply and demand curves can change rapidly.

We construct a large 2-dimensional test space of markets characterised by the slope of the supply and demand curves. We optimise Qs for each of these markets using a modified version of our ES and try to observe any regularities that emerge from this 3-Dimensional plot of optimal Qs values. We try to provide qualitative explanations for the reasons that might be behind the emergence of these observed trends: a full quantitative analysis is made difficult because of the presence of the order-queue as well as the multi-player nature of the CDA.

We fully realise that ultimately these explanations only proposed hypotheses at best. More analysis is needed to further investigate the validity of these hypotheses. We feel that such an effort will not only lead to a better understanding of the market mechanism that we have explored (using Qs) but also help gain deeper insight into the working of the CDA as an institution in general.

Chapter 5

Conclusion and Directions for Further Research

In this chapter we summarise our work and present directions for future research.
We highlight the implications of our results and argue that our work has presented, for the first time ever, very strong evidence for using an asymmetric auction mechanism for online auction markets regardless of the trading-strategy used.
We fully realise the difficulties that will be faced in determining the optimum Qs value for conducting an auction in a real-time dynamic auction market. We point out that further research is needed to help us determine the relationship between the underlying supply and demand curves in a given market and the Qs value that can be used to conduct an auction that will lead to larger gains than are possible with conventional market mechanisms.

Empirical economics has made rapid strides as a discipline since Smith organised his experiments with human subjects to study equilibrium formation in the CDA. Advances in computing technology and better modelling techniques have meant that we can now construct detailed models of complex social institutions like markets and study them in their true complex form.

The CDA, even though hugely popular, is also little understood, and gaining insight into its properties is important both from an economic as well as scientific perspective. CDAs are not only the most common type of auction mechanism in the world but they also find application in many artificial intelligence problems using the paradigm of Market Based Control (MBC). Some recent research has shown that the CDA, although difficult to analyse rigorously, lends itself to an experiment-based empirical approach. Our work can be looked upon as a step in this direction.

Although trading-agent design for the CDA is an established field and recently trading agents have been shown to outperform human traders, little work has been done in the field of automatic market design. This vista has become even more important of late because of the development of even better agents than those that were shown to outperform human subjects. It is foreseen that in the near future these agents could play a significant role in actual economic contexts.

Cliff had proposed an idea about market design in which the bidding process could be regulated by a parameter Qs – our work is an attempt to investigate this idea thoroughly and determine whether it has potential for wide application in actual marketplaces. We have used a different trading strategy (ZI-C agents) as well as a different performance measure and shown that hybrid variants of the CDA that are designed by an ES are superior to conventional mechanisms for many of our test cases. These are the first-ever results that demonstrate that hybrid markets can dominate conventional variants with more than on type of trading algorithm. More research is needed to prove that these results hold for all types of trading strategies and perhaps even human traders. More research is also needed to help us determine exactly how this parameter leads to gains in efficiency in different market contexts. Also we note that outcomes for experiments performed for the same markets but with different trading strategies can be different. It appears that the optimal Qs value for a given market is dependant upon the trading strategy used. Experiments with large markets populated with a large number of heterogeneous trading-agents could help us show if this indeed the case and whether hybrid variants emerge as optimal even in these markets.

We recognise that a full potential of this idea for automatic market design can only be realised if we gain better understanding of how it works, and are able to devise faster and more accurate ways of designing optimal markets with minimal risk in situations where we have only incomplete or inaccurate information. We have tried to do this by empirically studying patterns which can emerge from conducting a large number of tests over a search space of test-markets characterised by given parameters. We have shown that certain intuitions are confirmed by our experiments and the results agree with some previously known fact about the design of conventional market-mechanisms. Analysis of multi-player games like markets is hard. We still need more insight into the relationship of the optimal Qs value with underlying market characteristics. Such research can help us in applying these ideas to real-time 'live' markets and ultimately lead to better market dynamics that can help us maximise social welfare outcomes.

APPENDIX

Appendix A

Using Agent-Based Simulations to Study Complex Economic and Social Systems: A Review

In this appendix we motivate the case for agent-based simulation of economic and social systems. We highlight the shortcomings of the traditional tools that have been used for modelling these systems and argue that an agent-based, bottom-up approach allows greater flexibility and insight into the functioning of these systems.

Economics can be defined as the study of those phenomena that can be understood as emerging from the interactions among intelligent, self-interested individuals (Krugman, 1996). It is interesting to observe that this definition of economics is remarkably similar to that of a typical evolutionary system. It gives rise to the interpretation of a modern day market-based economy as a complex adaptive system in which a decentralised collection of intelligent agents interact with each other in various market contexts.

Even though this interaction happens at a local level in a distributed and parallel fashion, it gives rise to global regularities like equilibrium formation and high market efficiency in trading marketplaces and the emergence of trade-networks, socially accepted monies, market protocols, business cycles and the common adoption of technological innovation. It is important to note that there is no a-priori incentive for the emergence of any of these regularities – they emerge solely out of the low-level interactions between the agents of the system.

These macro-level regularities then feed back into the system, affecting the determination of the micro-level local interactions. The result is a complicated dynamic system of recurrent causal chains connecting agent behaviour, interaction networks and social welfare outcomes.

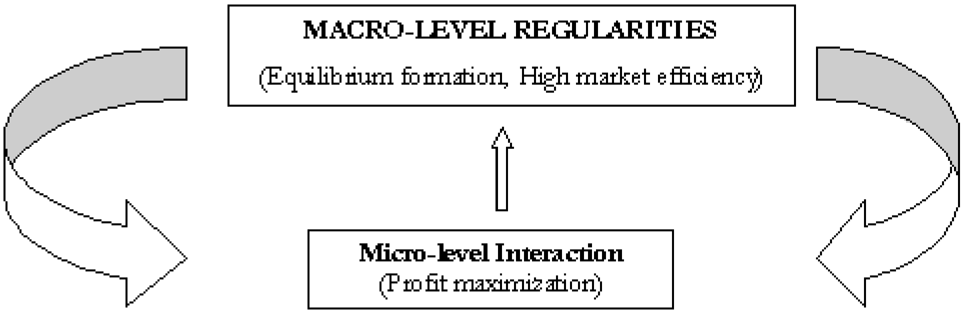

Figure A.1: Two-way feedback between microstructure and macrostructure

Building Economies Top Down

Even though this two-way feedback between microstructure and macrostructure has been recognized within economics for very long (Smith, A. 1937; Olson, M. 1965; Schelling, T. 1978), traditionally, it is the global regularities rather than low-level agent interactions that have been studied extensively in economics. The reasons for concentrating on the macro-level properties are obvious. It is difficult, if not impossible, to model and analyse the large number of individual micro-level interactions between the agents that make up the system.

On a typical trading day the New York Stock Exchange (NYSE) witnesses more than 20 million trades with financial transactions valued at over 33 billion dollars[15]! To record and analyse such a large volume of trades would be a non-trivial exercise. Even so – it would not model several of the micro-level properties of the system like price negotiation, unsuccessful bids and agent-agent interaction, all of which are important aspects of the low-level agent behaviour.

[15] Source: NYSE official statistics for 20th of August 2002 from NYSE's official website http://www.nyse.com.

In order to simplify task of modelling a system that is inherently complex, economists often make certain assumptions about the behaviour of agents that are active in the given system (Krugman, 1996). Chief of these are that: -

a) The agents are not only intelligent, but they *maximise* i.e. they choose the best of all feasible alternatives and
b) They achieve an *equilibrium*, in which each individual is doing the best she can, given the market conditions.

Conventional models built on these assumptions are extremely elegant in form. However, it is easy to see that these assumptions are extreme and unrealistic. Not only do they contradict everyday experiences (all of have been 'duped' into buying goods for more than the lowest price in the market), models built based on these assumptions have failed to explain even the most basic empirical features of real markets (Tesfatsion, L. 2002).

Conventional economic models are built top-down. Heavy reliance is placed on externally imposed coordination devices such as fixed decision rules, knowledge assumptions, representative agents and market equilibrium constraints. Low-level interaction between agents plays little or no role in these models for the simple reason that till date, economists have simply lacked the means to model it in anywhere near its true complexity.

Building Economies Bottom Up

In her exhaustive survey titled "Agent Based Computational Economics: Growing Economies from the Bottom Up", Leigh Tesfatsion (Tesfatsion, L. 2002) convincingly argues that a paradigm shift is happening in the way economic and social systems are modelled. This paradigm shift has been made possible by the advent of powerful new computational tools, in particular the development of new techniques in evolutionary computation, which gives economists an opportunity to model market dynamics in their true complex form. There is a rapidly growing body of research in the field (Tesfatsion 2001a; 2001b; 2001c) in which the Agent-Based approach has been applied to finance, electricity auction and even retail-coffee markets[16]. This research has led to extremely encouraging results with the models providing insight into, and possible explanations for a variety of observed regularities in the markets.

Bunn and Oliveira (2001) report a study in which the proposed New Electricity Trading Arrangement (NETA) for the UK was modelled using a multi-agent evolutionary model. This agent-based model was able to provide pricing and strategy insights, ahead of the NETA's actual introduction in November 2001. Nicolaisen, Petrov and Tesfatsion (2001) also report a study in which the performance of an electricity market with a new auctioning mechanism is simulated using agents using reinforcement learning. The study provides key insights into the sensitivity of market efficiency to the learning behaviour of the agents when the proposed market-mechanism is used. The study concludes that the proposed auctioning mechanism is not robust against the active exercise of bad judgement on part of the agents. Such research can act as a validation tool for proposed market structures without incurring heavy costs.

Modelling Social Structures

Carley (2001) also presents a survey of the field in which she highlights the application of agent-based modelling techniques for studying modern-day social organisations. She presents the view that any entity composed of intelligent, adaptive and computational agents is also an intelligent, adaptive and computational agent. This property of an organisation is termed as *synthetic adaptation*. She puts forwards two important reasons why agent-based simulations of social systems will become integral tools for those wishing to gain insights into the dynamics of socio-technical systems such as organisations.

Firstly, organisations are inherently complex. The capabilities and overall properties of any organisation emerge from the detailed low-level interactions and behaviour of the member agents. However, this emergent behaviour is not the result of a simple

---

[16] Leigh Tesfatsion maintains an Agent Based Computational Economics (ACE) website at http://www.econ.iastate.edu/tesfatsi/ace.htm.

The site provides extensive resources to those interested in the field of ACE research aggregation across member behaviour but inherently complex and non-linear in nature. Agent-based simulations are the best way to gain insight into this complex process.

Secondly, the nature of organisations is increasingly changing from those trading in tangible goods to those that trade in information. This has challenged many of the premises that traditional organisational theory was based on. Organisations need to optimise their structures in order to meet the challenges presented by the new information economy. They need to reduce their reaction times and make fast decisions based on inaccurate or incomplete information. In order to do this, the organisation needs to move through the structure search-space searching for a form that enables higher performance (Levinthal et al 1981). However, structural changes in organisations have high costs and potentially devastating consequences so any changes must be initiated only after careful deliberation. Agent-based simulations allow us to answer "what-if" questions about the structure of these organisations aiding us to understand how adoption of different technologies, decisions and policies influence the performance, effectiveness, flexibility and survivability of an organisation. These models enable us to run large-scale experiments without having to alter the structure of the organisation in any way and are invaluable tools for policy analysis.

It is also worth mentioning the research in the field of social behaviour science. Robert Axelrod's 1984 book "The Evolution of Cooperation" was enormously influential among economists, game-theorists and social scientists alike. Axelrod (1984) demonstrated that mutual co-operation could evolve among self-interested, non-related agents through reciprocity, with little or no explicit forward-looking behaviour on part of these agents. Various other emergent social phenomenons like segregation by race (Schelling, 1978, P. 147-155) and the emergence and stability of equity norms in society (Epstein et al 1996) using simple boundedly rational agents have been studied by social scientists. These works have supported the use of simple boundedly rational agents for the study of social phenomenon.

It is clear that humans and the systems in which they are embedded are not perfectly-rational in the economic-sense. They *satisfice* rather than maximise: i.e. they are cognitively limited. Hence, it is difficult for us to choose the level of sophistication of cognitive-behaviour that the agents in our simulations exhibit. In general choosing the level of sophistication or rationality can be a difficult design decision and it can affect the richness of the simulation significantly.

Summary

Economies and social structures can be imagined as being complex adaptive systems in which a decentralised collection of intelligent agents interact with each other in contexts. Even though this interaction happens at a local level in a distributed and parallel fashion, it gives rise to global regularities and it is the global regularities rather than low-level agent interactions that have been studied extensively in economics and the social sciences.

However a paradigm shift is happening in the way economic and social systems are modelled today and increasingly agent-based simulations are being used to study diverse structures like finance, electricity auction and retail-coffee markets, organisation structures and emergent social phenomenon observed in society. This research has led to extremely encouraging results with the models providing insight into, and possible explanations for a variety of observed regularities.

Agent-based simulations allow us to answer "what-if" questions about the structure of these organisations aiding us to understand the affect of factors of interest on these systems in a controlled and tightly regulated environment. A wider adoption of these methods will allow us to gain insight into systems that have remained mysteries so far due to a lack of tools to study them.

Appendix B

Running the Code

The code has been stored in the folder:
`/home/students/msc/msc80vxw/Summer_Project`

The code has been written in C++ and is compiled using the gnu C/C++ compiler `gcc` (`gcc` can be downloaded free from the GCC homepage http://gcc.gnu.org). Two shell files called `compile.sh` and `fastcompile.sh` are stored in the folder and can be used to compile all programs together. `compile.sh` generates binaries with debug information (for `gdb` or compatible debuggers) and `fastcompile.sh` without debug information. Binaries without debug information can be faster to execute.

The code is properly indented and commented and should be easy to read, understand and modify. All input and output files are simple ASCII test files. We outline the most important parameters and how to change their values below.

General parameters, which affect most programs, are stored in `MarketParams.h`. These include: -
`MARKET_POP`: Size of population for ES
`NO_OF_GENERATIONS`: Maximum number of generations the ES will run for
`REPEAT_TRIALS`: The number of trials over which any fitness evaluation is averaged

Now we discuss the specific programs.

EA

Contains the code for the ES used for experiments in Chapter 4. The executable is called `ea` and is run from the command line as follows:

```
Usage:ea buyerFile sellerFile QsLogFile FitnessLogFile
StepSizeLogFile
```

Where `buyerFile` and `sellerFile` are text files containing the limit prices for a given market. These simple lists of numeric vales separated by a white space (space, tab or newline) and are of the form:

```
8
7
1
7
2
7
```

The `QsLogFile` and `FitnessLogFile` logs are stored with a (Gen, Value) format. `StepSizeLogFile` lists the step sizes for each successful mutation during the run.

QsTester

This is the landscape generation utility. It is run as follows:

```
Usage:QsTester buyerFile sellerFile QsLandscapeLog StdDeviationLog
```

`QsLandscapeLog` contains the fitness value for the landscaped values of Qs in a (Qs, Fitness) format. The `StdDeviationLog` lists the Standard deviation values for each of these points in a (Qs, StdDeviation) format.

The parameters for this program are: -

`REPEAT_EXPERIMENT`: allows us to generate landscapes incrementally. This way we can set `REPEAT_TRIALS` to a moderate value (1000) and repeat the landscaping process many times (say 10) and generate exhaustive landscapes (in this case with fitness values averaged over 10,000 repetitions) with a gradual increase in accuracy. The program reports when a repeat is complete and the output files can be used to check the nature of the landscape at various stages in this process.

`QS_POINTS:` The number of Qs values for which the landscape is plotted. The points are equally distributed in the [0,1] interval.

QsMap

This is the program used to generate 3-D landscapes of the form (mB,mS,optQs) that was used to generate the results discussed in Chapter 5. The executable is called `QsMap` and is used as follows:

```
Usage:QsMap QsMapLogFile
```

`QsMapLogFile` stores the output in a (mB,mS,optQs) format. The parameters for this program are contained in the file `QsMap.h`. Some of these are:

`MIN_LIMIT`: The lowest price in the market. This value is referred to as minPrice in Chapter 5.

`MAX_LIMIT:` The highest price in the market. This value is referred to as maxPrice in Chapter 5.

`MIN_SLOPE:` The lowest value mB or mS can take.

`MAX_SLOPE:` The highest value mB or mS take.

`SLOPE_INCREMENT:` The step size between successive mB or mS values.

`MAX_STAGNANT:` Discussed in detail in Chapter 5.

Bibliography

Adam Smith *The Wealth of Nations*, The University of Chicago Press, Chicago, IL, 1776

Arthur W.B. "Self-Reinforcing Mechanisms in Economics" In *The Economy as an Evolving Complex System*, ed. Anderson P., Arrow K., and Pines D., 9-32. Santa Fe Institute Studies in the Sciences of Complexity, Proceedings Volume 5, Redwood City, CA. Addison-Wesley, 1988.

Axelrod, R. *The Evolution of Cooperation*, New York, NY: Basic Books, 1984

Back T. and Schewefel H. P. " An overview of evolutionary algorithms for parameter optimisation", *Evolutionary Computation*, Vol. 1, No. 1, p. 1-23, 1993

Begg, D., Fischer, S. and Dornbusch R. *Economics (Fourth Edition)*, McGrawHill, London, 1994

Becker, Gary S. "Irrational Behaviour and Economic Theory." *Journal of Political Economy*, 70, 1-13, February 1962

Blake LeBaron "Agent based computational finance: Suggested readings and early research", *Journal of Economic Dynamics and Control,* 24, 679-702, 2000

Bollerslev, T., Chou, R. Y., Jayaraman, N. and Kroner, K. F., "ARCH modelling in finance: A review of the theory and empirical evidence", *Journal of Econometrics*, 26, 318-319, 1990

Chris Preist and Maarten van Tol "Adaptive agents in a persistent double shout auction". *In Proceedings of ICE-98.* ACM 1998

Clearwater, Scott Ed. *Market-Based Control: A Paradigm for Distributed Resource Allocation*, World Scientific, 1995.

Cliff D. "Evolutionary Optimisation of parameters sets for adaptive software-agent traders in continuous double-auction markets". Technical Report HPL-2001-99, Hewlett-Packard Laboratories, 2001

Cliff D. "Evolution of market mechanism through a continuous space of auction types". Technical Report HPL-2001-326, HP Labs, 2001a

Cliff D. "Evolution of market mechanism through a continuous space of auction-types II: Two-sided auction mechanisms evolve in response to market shocks", Technical Report HPL-2002-128, HP Labs, 2002

Cliff, Dave and Janet Bruten, "Simple Bargaining Agents for Decentralised Market-Based Control", Technical Report no. HPL-98-17, Hewlett-Packard Laboratories, 1998

Cliff, Dave and Janet Bruten "Zero is not enough: On the lower limit of agent intelligence for continuous double auction markets", Technical Report no. HPL-97-141, Hewlett-Packard Laboratories, 1997

Das R., Hanson J. E., Kephart J. O. and Tesauro G. "Agent Human Interaction in a Continuous Double Auction", In *Proceedings of the International Joint Conferences on Artificial Intelligence*. Seattle, USA, August 2001

Dave Cliff and Chris Priest "Days without end: On the stability of experimental single-period continuous double auction markets", HP Labs Technical Report, HPL-2001-325, Dec 2001

Derek Bunn and Fernando Oliveira "Agent-Based Simulation: An Application to the New Electricity Trading Arrangements of England and Wales", *IEEE Transactions on Evolutionary Computation*, Volume 5, Number 5, 493-503, October 2001,

Epstein, J.M. and Axtell, R. *Growing Artificial Societies: Social science from the Bottom Up* Cambridge, MA: MIT Press, 1996

Fogel D. B. *Evolutionary Computation: Towards a New Philosophy of Machine Intelligence*, New York, NY: IEEE Press, 1995.

Fogel D. B. *Evolving Artificial Intelligence*, PhD thesis, University of California, San Diego, CA, 1992

Fogel D. B. *System Identification Through Simulated Evolution: A Machine Learning Approach to Modelling*, Ginn Press, Needham Heights, MA 02194, 1991

Fogel L. J., Owens A. J. and Walsh M. J. *Artificial Intelligence through Simulated Evolution* John Wiley and Son, New York, NY, 1966

Friedman D. "The double auction market institution: A survey", In *The Double Auction Market: Institutions, Theories, and Evidence*, D. Friedman and J. Rust, Eds, P.. 3-25, Addison Wesley, 1993.

Gerald Tesauro, Rajarshi Das "High-Performance Bidding Agents for the Continuous Double Auction", *Proceedings of the ACM Conference on Electronic Commerce (EC-01)*, Tampa, Florida, October 14-17, 2001.

Gode D.K. and Sunder S. " Lower Bounds for Efficiency of Surplus Extraction in Double Auctions", In *The Double Auction Market: Institutions, Theories, and Evidence*, D. Friedman and J. Rust Eds, Addison Wesley, 1993.

Gode D.K., and Sunder, S. "Allocative Efficiency of markets with Zero-Intelligence Traders: Markets as a partial substitute for individual rationality" *Journal of Political Economy*, 101, 119-137, 1993

Goldberg D. E. *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison-Wesley, Reading, MA, 1989

Hansen, L. and Singleton, K. "Stochastic consumption, risk-aversion, and the temporal behaviour of asset returns", *Journal of Political Economy*, Volume 91, 1983

Holland J. H. *Adaptation in Natural and Artificial Systems*, The MIT Press, Cambridge, MA, 1992

James Nicolaisen, Valentin Petrov, and Leigh Tesfatsion "Market Power and Efficiency in a Computational Electricity Market with Discriminatory Double-Auction Pricing", *IEEE Transactions on Evolutionary Computation*, Volume 5, Number 5, 546-560, October 2001

Kathleen M. Carley forthcoming, "Computational Organization Science: A New Frontier" in Arthur M. Sackler Colloquium Series on Adaptive Agents, Intelligence and Emergent Human Organization: Capturing Complexity Through Agent-Based Modelling October 4-6, 2001, Irvine, CA, National Academy of Sciences Press.

Kephart J. "Software agents and the route to the information economy", *Proceedings of the National Academy of Sciences of the United States of America*, 99, No. Suppl, 7207-7213, 2002.

Krugman P. "What economists can learn from evolutionary economists", Talk given to The European Association for Evolutionary Political Economy, November 1996, http://www.mit.edu/~krugman/evolute.html

Levintal, Daniel A. and James G. "A Model of Adaptive Organisational Search", *Journal of Economic Behaviour and Organisation*, Volume 2, 307-333, March 1981.

M. Matsumoto and T. Nishimura. "Mersenne twister: A 623-dimensionally equidistributed uniform pseudorandom number generator", *ACM Transactions on Modeling and Computer Simulation*, 8(1), 3-30, 1998

M. Matsumoto "The mersenne twister home page" http://www.math.keio.ac.jp/matumoto/emt.html, December 1997

Michael P. Wellman "The economic approach to artificial intelligence" *ACM Computing Surveys*, 27(3), 340-342, September 1995.

Olson, M. 1965. *The Logic of Collective Action.* Harvard University Press.

Phelps S., McBurney P., Parsons S., Sklar E. "Co-evolutionary auction mechanism design: a Preliminary Report", Technical Report ULCS-02-019, Department of Computer Science, University of Liverpool, 2002
http://www.csc.liv.ac.uk/research/techreports/tr2002/ulcs-02-019.ps

Preist C. "Economic Agents for Automated Trading", In *Software Agents for Future Communication Systems*. Hayzelden A.L.G. and Bigham J. Eds, Springer, 1999

Press, W.H., Flannery, B.P., Teukolsky, S.A. and Vetterling, W.T. *Numerical Recipes in C*, Second Edition. New York: Cambridge, 1992

Richard M. Stallman *GNU C Compiler Manual Version 3.1.1*, The Free Software Foundation, Boston, MA, 2002

Rust J., Miller J. and Palmer R. "Behaviour of trading automata in a computerised double auction market" In *The Double Auction Market: Institutions, Theories, and Evidence*, D. Friedman and J. Rust, Eds., Addison Wesley, 1993.

Samuelson, P. A. *Foundations of Economic Analysis*. Cambridge, Mass.: Harvard University Press, 1947.

Schelling,T. *Micromotives and Macrobehaviour*, W.W. Norton, New York (1978).

Schwefel Hans Paul *Evolution and Optimum Seeking*, John Wiley and Sons, New York, 1995

Schwefel, Hans Paul *Numerical Optimisation of Computer Models*, John Wiley and Sons, Chichester, 1981

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

Smith, A. *An Inquiry into the nature and causes of the wealth of nations*, New York, NY: Cannan Edition, American Modern Library Series, 1937.

Smith, V.L. "Microeconomic Systems as an Experimental Science", *American Economic Review*, 72, 923-955, 1982

Sweeney, R. J., "Beating the foreign exchange market", *Journal of Finance*, 41, 163-182, 1986

Tesfatsion, L. Guest Editor, Special Issue on Agent-Based Computational Economics, *Computational Economics*, Vol. 18, No. 1, October 2001b

Tesfatsion, L. Guest Editor, Special Issue on Agent-Based Computational Economics, *Journal of Economic Dynamics and Control*, Vol. 25, No 3-4, March 2001a

Tesfatsion, L. Guest Editor, Special Issue on Agent-Based Modelling of Evolutionary Economic Systems, *IEEE Transactions on Evolutionary Computation*, Vol. 5, No. 5, October 2001c

Tesfatsion, Leigh "Agent based Computational Economics: Growing Economies from Bottom Up", *Artificial Life*, Vol. 8, No 1, 55-82, 2002

Varian H. R. "Economic mechanism design for computerized agents", In USENIX Workshop on Electronic Commerce, New York, NY, July 1995

Yao X., Liu Y. and Lin G. "Evolutionary programming made faster," *IEEE Transactions on Evolutionary Computation*, **3**(2):82-102, July 1999