Real Time Trading Mechanisms for Automated Markets

Rehan Abbas Zaidi

Department of Electronic Engineering
Queen Mary, University of London

Submitted to the University of London in partial fulfilment of the requirement for the degree of
Master of Philosophy

Supervisor: Dr. Chris Phillips

December 2004

Abstract

The growth of electronic commerce has aroused considerable interest in the investigation of mechanisms for trading commodities online. The potential for trading institutions to avail of automated trading environments and trading agents has spawned wide-ranging research into negotiation and interaction protocols in multi agent systems. The theory of mechanism design from economics has been employed to construct entities that allow autonomous, self-interested software agents to interact in distributed environments.

This work provides a critical survey of the field of automated markets at the confluence of microeconomic theory and multi-agent systems with particular emphasis on market mechanisms that provide a continuous trading environment. The long-term objective of the research in this field is the creation of fair and accessible market environments for personalised agents to operate within. The work presented here analyses this objective from several perspectives, the system designer providing the protocols/regulations, the preferences of consumers trading in the market and the service providers trying to balance obligations with opportunities. It includes analysis of trading models, implementation of market mechanisms, development of algorithms and protocols for automatic negotiation and bidding strategies for agents.

An in-depth survey of automated commodity markets is undertaken with particular emphasis on double auctions. The continuous double auction mechanism is investigated through simulation studies, which demonstrate the need for a high degree of synchronization between buyers and sellers for optimal performance in terms of efficiency and liquidity. However, such synchronization is unlikely to occur in practice in a continuous real-time environment, so there is a need for a market mechanism to provide sufficient liquidity and immediacy. In this context, the thesis considers a new type of automated market, i.e. a quote-driven market, where market maker(s) aggregate demand and supply continuously over time. Quote-driven markets are suitable for real-time trading in commodities e.g. financial products, bandwidth, and electricity. This work outlines typical scenarios for emerging commodity markets in electricity and bandwidth trading within which the operation of various market mechanisms is assessed.

Acknowledgements

This work was partially funded by the EPSRC under grant GR/R10394/01.

I would like to express my gratitude to *Agent Link* for the support provided for me to attend ACAI-01 and the Third European Agent Systems Summer School.

I am also indebted to several people who have helped me during the course of this research in one way or another. In particular I would like to thank the following:

- *Dr. Rachel Bourne* for providing me the opportunity to take up this research and for her supervision and guidance during the initial course of this work.

- *Dr. Chris Phillips* for his supervision during the latter part of this work and above all his constant encouragement and support without which I would not have progressed.

- The *staff* at the Electronic Engineering Department, QMUL for their support and facilitation of this research and making my stay welcome.

- *Janette Paterson* and *Peter Foulkes*, project managers on the NETA project, LogicaCMG for their consideration and support during the course of writing up the thesis.

- *Prof. Vernon Smith, Prof. Dan Gode and Chris Priest* for their clarification and advice when I sought it.

- My uncle *Muzaffar Mohamed* and aunt *Nikhat F. Mohamed* who were instrumental in my taking up studies in the UK and who have supported me throughout.

- Last but not the least I would like to thank my family for their love and encouragement, and for their belief in me which serves as the guiding light in my quest.

Dedicated to the loving memory of my grandmother

(Late) Mrs. Naseem Fatima

List of Acronyms

3G	Third Generation
AI	Artificial Intelligence
B2B	Business to Business
BETTA	British Electricity Trading and Transmission Arrangements
CAC	Connection Admission Control
CAPM	Capital Asset Pricing Model
CBOT	Chicago Board of Trade
CDA	Continuous Double Auction
CDPS	Cooperative Distributed Problem Solving
DA	Double Auction (used interchangeably with CDA)
DAI	Distributed Artificial Intelligence
FX	Foreign Exchange
GUI	Graphical User Interface
ISO	Independent System Operator
MAS	Multi Agent System
MDP	Market Data Provider
MVNO	Mobile Virtual Network Operators
NSC	National Stock Exchange (U.S.)
OFCOM	Office of Communications
OFGEM	Office of Gas and Electricity Markets
OFTEL	Office of Telecommunications (superseded by OFCOM)
QD	Quote Driven
QOS	Quality of Service
SEBI	Securities and Exchange Board of India

SETS	(London) Stock Exchange's Electronic Trading Service
SLA	Service Level Agreement
SP	Service Provider
S/W	Software
NASDAQ	National Association of Securities Dealers Automated Quotation System
NETA	New Electricity Trading Agreement
NP	Network Provider
NYSE	New York Stock Exchange
XETRA	Exchange Electronic Trading (trading platform of Deutsche Börse Group)
ZI	Zero Intelligence
ZI-C	Zero Intelligence – Constrained
ZI-U	Zero Intelligence – Unconstrained
ZIP	Zero Intelligence Plus

Chapter 1
Introduction

Two distinct but impacting trends that have come to characterise the networked world of today are the increasing levels of decentralisation and automation in computing. The future could see an open free market economy of software agents [IBM-Research] interacting in a variety of ways in pursuance of specific goals or tasks on behalf of their users. In such a scenario it is entirely plausible for us to turn to economics for a solution to satisfying the often competing and divergent needs of the self-interested software agents interacting therein. The concept of a free market economy advocates a good or optimal allocation of resources arising out of the actions of individuals engaged in trading conventions driven purely by self-interest. In such situations, the notion of a self-interested agent as in AI terminology corresponds closely with that of a rational, utility maximising agent in microeconomic theory. Economists have studied the relationship between market structures and efficiency. The *market* acts as an institution which concentrates buyers and sellers and defines the rules for interaction i.e. the *protocols*. There are several market structures studied in literature, which implement a wide variety of mechanisms in respect of the actual trading conventions. However they all have a central theme running through them, namely that of the market mechanism (market structure) being able to compute a *market clearing price* or *equilibrium price* where the quantity demanded matches the quantity supplied. This price determination is achieved as a result of (price) competition between the agents in the market. Therefore at this competitive equilibrium, a commodity will have been efficiently allocated amongst the competing agents.

This chapter presents the necessary background and the motivation for the research. It lists in general, the contributions of this research and presents an overview of the following chapters.

1.1 Agents, Interactions and Self-Interest

In the not too distant future, the Internet could be populated by a multitude of *self-interested* intelligent agents pursuing goals, performing certain tasks on behalf of their users or providing services to other agents/humans[1]. These software entities or agents will increasingly represent human beings in tasks ranging from information retrieval and dissemination to trading scenarios involving complex negotiation mechanisms. This potential for software agents acting on behalf of individuals and businesses in real world commercial applications has stimulated a rapid increase in research into negotiation and interaction protocols in multi-agent systems. The ultimate goal of multi agent systems research is undoubtedly to develop mechanisms that enable agents to interact as well as humans. Having said this, it is unquestionably a very demanding task to construct mechanisms that are capable of capturing complex negotiations as richly as they occur in the real world. Recent work has focused on defining these mechanisms in more restrictive settings. One large class of these is represented by economically motivated interactions e.g. buying and selling of goods etc. This introduces a great deal of simplicity in defining the protocols because if the product being bought or sold is 'standardised' in some way (perhaps by means of benchmarks etc) the only issue that needs discussion is the price of the '*commodity*'. Therefore the interactions between agents can mostly be couched in terms of price i.e. bids and/or offers.

Sandholm [Sandholm 1999] attributes the importance of negotiation systems in a society of self-interested agents in part to the need for distinct agents (belonging to different organizations) to interact in an open environment over an ever-growing standardized communication infrastructure. Another reason cited is the advantages offered by computational agents for negotiation at the operative decision making level. Besides saving labour for humans, agents can be more effective at obtaining beneficial results in strategically and combinatorially complex environments. Yet another important contributor is the development of a decentralised computing infrastructure. The trend towards decentralised computing has led to a completely new paradigm of distributed open networks. This has prompted people to develop new metaphors for thinking about the future of computing. Distributed systems can be thought of as open and self-regulating entities acting as individuals or agents. Huberman [Huberman 1988] describes such systems as showing many characteristics of social or biological organizations[2]. In the context of such a distributed system, the behaviour of the entities

---

1 See the Semantic Web (http://www.w3.org/2001/sw) and the Agentcities initiative (http://www.agentcities.org) for a perspective into the future agent based networked world.

2 IBM have developed on this notion in their blueprint for autonomic computing http://www-3.ibm.com/autonomic

(humans or agents) in their interactions, strategies and competition for resources models whole ecologies[3]. Agent technology got its initial impetus from the use of agents as an abstraction tool in the design of such systems.

In a scenario where agents are not cooperative but self-interested with private information and goals, the crucial problem is one of designing *incentive compatible* protocols, which compute optimal system wide solutions despite the self-interest of the individual agents. The system designer in such a case does not control the behaviour of all the system components. S/he can only control the *mechanism* or *protocol* (rules of the game/interaction) while the individual agents choose their *strategies*. With incentive compatible protocols, the system designer would like to make sure that the agents are motivated to behave in a desired manner, which ensures a desirable overall outcome. Economics offers us valuable insights into implementing coordination mechanisms in computational systems [Roth 2002], as it is essentially a science that deals with allocation of scarce resources in human society. Similar considerations arise in distributed interconnected networks where sharing of scarce resources efficiently (be it processor time, network storage or bandwidth etc.) is of utmost importance. In economics, a mechanism of central importance in human resource allocation and decision-making is the market. The market allows agents to express demands and allocate goods and services in response to these demands. The notion of market mechanisms as a tool for negotiation in multi agent systems is convenient because it solves two potential problems to a large extent. These are: -

1) The ontology problem i.e. the market protocol be it an auction mechanism or some other protocol, provides a common vocabulary or ontology which characterizes the agent interaction in such domains.
2) The problem of searching for suitable/favourable parties. The market institution serves to provide a focal point, i.e. an aggregator for interested parties (i.e. search costs are reduced).

---

3 An interesting albeit parallel field of research, which studies the complex interactions and outcomes in Multi Agent Systems in a much broader scope, is Complexity (see Lamper, D. and N. F. Johnson (2002). "The Science of Complexity: Finding out when two's company and three's a crowd." Dr. Dobb's journal Intelligent Systems 341: 16-22)

The above factors have led to the development of an emerging multidisciplinary research area, which is at the convergence of AI, economics, game theory and algorithmic theory.

1.2 Automated Markets

It will be useful at this stage to classify the broad areas of work at this interface of AI, economics, game theory and algorithmic theory. The term c*omputational markets,* is generally used to denote a market in which computer programs or computational agents trade commodities with each other. Within this larger field, Ygge [Ygge 1998] categorizes the work as under :-

1) Simulation of markets
2) Resource allocation implementations.

The first category is essentially a laboratory study of market environments with computer programs replacing human traders. Economists find this useful to develop methods that describe and formalize human behaviour. Essentially this research aims to provide new insights into economic theory.

The second category of resource allocation implementations use the abstractions developed in the first instance above to design and implement systems. The term *Market-oriented programming* [Wellman 1996] is sometimes used to describe implementations of resource allocation mechanisms in computer systems. These can be further divided into two sub-categories: -

a) Implementation of real markets
b) Resource allocation/optimisation problems where the market is used for maximising (or minimising) some global measure. It is seen in markets for network bandwidth, memory/processor allocation, temperature control and information searching/selling etc.

This thesis is concerned with the first aspect i.e. automating participation in electronic markets, although there will be instances where the second issue will have a bearing. Through automated trading by software agents, besides obtaining improvements in the quality of existing markets, such as consumer goods markets, service markets, and the emerging information markets, we can reap the benefits of markets as effective instruments of resource allocation in non-traditional domains too, such as fine-grained markets for electric power and communication bandwidth. Agents have the capacity to consider more information, e.g., evaluate thousands of offers for a new car and hundreds of recommendations from various sources, and may also act in domains where we are disqualified due to speed requirements.

1.3 Auctions, MAS and Electronic Markets

The traditional model of pricing goods has followed the fixed pricing structure since the industrial age when mass production of consumer goods began. However with the emergence of the Internet and electronic commerce this is being challenged by the dynamic pricing model. In dynamic pricing systems the price of the commodity is determined continuously by the expression of demand and supply by the participants. Essentially, as this model sees a continuous fluctuation in prices, there is a requirement for aggregating the demand and supply in the market to effectively establish the price for the goods. Auctions have traditionally served as the mechanism for performing this aggregation in the real world. An auction can be thought of as a disinterested mediator, which simply follows a formal policy that defines its behaviour as a function of the bids it receives. The use of auctions as a dynamic pricing system has long been established in markets for items like securities, airline tickets and oil. Recently the growth in e-commerce can also be attributed in large measure to the success of employing auctions as a mode for business exchanges both in the Business-to-Consumer and Business-to-Business domains[4]. Auctions offer numerous advantages [Kersten and Lo 2001] including process efficiency, ease of use, small transaction costs, reach and ability to manage a large number of bidders as well as their ability to manage the uncertainty and ambiguity of value in a social context.

Auctions are a convenient mode for mechanism design in multi agent systems because they have provable properties and are often readily analysable [Sandholm 1999]. Given that auctions are a convenient means for resolving one to many (or many to many) negotiation situations, they have been used as internal resource allocation devices in Multi Agent Systems or MAS [Clearwater 1996; Gibney, Jennings et al.

---

[4] The Internet Auction List website www.internetauctionlist.com currently provides more than 2500 auction company listings.
Juniper Communication predicts that B2B e-marketplace transactions will rise to $137 billion by 2005 while Forrester predicts that auctions will account for 25% or $54 billion of online retail sales in the US alone by 2007.

1999; Ygge and Akkermans 1999]. Indeed, if time is discounted or if communications are costly, auctions offer an effective means as a negotiation device in MAS.

1.4 Continuous Trading Environments and the need for "immediacy"

As mentioned above, auctions aggregate the demand and supply in a market to determine the clearing price. However this implies a high degree of synchronization between the buyers and sellers for the auction to establish an efficient outcome (i.e. a competitive equilibrium). In a continuous trading environment, for instance in financial markets, customer orders for execution arrive randomly and spread over time i.e. the demand and supply in the market is continuously fluctuating and it is difficult or impossible to establish an effective clearing price at any instant. In such a situation there exists a demand for immediacy providers who can service *buy* orders at a somewhat higher price and *sell* orders at a somewhat lower price. Several financial markets worldwide (including NASDAQ[5], NYSE[6] etc.) have such dealers (known as market-makers or specialists) whose presence serves to provide liquidity in the market. Such institutions (i.e. *Quote Driven* markets) provide an environment where intermediate market makers are obliged to quote live prices (two-way firm price or a *quote*) to clients, thus providing a willing counter-party to a trade at all times. In such a scenario the buyers and/or sellers in the market are willing to pay a premium (the difference between the *bid* and *offer* price), since they are guaranteed to be able to trade regardless of the activities of others. Quote Driven markets also serve to provide price stability by smoothing out fluctuations in the price level [SEBI 2000]. This is difficult in auction based (or *order driven*) markets, where price determination by balancing the demand and supply in the market may lead to huge price swings simply because of the natural intervals that intercede the arrival of orders.

1.5 Relevance of Research

The growth of electronic commerce involving the trade of goods and services online has resulted in considerable interest in the mechanisms for trading commodities. This together with the potential for autonomous software agents to act on behalf of

---

5 Securities market in the US, popular for its listing of major technology companies
6 New York Stock Exchange or the Dow Jones

individuals and businesses in real world commercial applications has spawned wide-ranging research into negotiation and interaction protocols in multi agent systems (e.g. the Information Economies Project at IBM[7] and the E-Services Markets Group at HP Labs[8]). The organisation of market institutions has become even more important in the context of electronic markets and the use of market mechanisms for negotiation and resource allocation/optimisation problems.

To be able to realise the true potential of automated markets we need a better understanding of such markets in order to create an accessible and fair environment where human beings are comfortable with the operation of their personalised agents. Such an exercise will involve work on several fronts, as one has to cater to the differing perspectives of the individuals involved e.g. the system designer would like to provide protocols and regulations that will enable the market to operate with minimum external interference. The consumers on the other hand would like to obtain reasonable prices, the ability to trade immediately etc. The service providers/market makers would like to have a balance between obligations and opportunities enabling them to make profits proportionate to their exposure to risk. The whole chain needed to create trusted electronic marketplaces include: *analysis of trading models, implementation of market mechanisms, development of algorithms and protocols for automatic negotiation mechanisms, development of bidding strategies for use by agents* etc. All these considerations beget the question: *What type of market environment can best achieve this situation and under what conditions*? This thesis takes a designer's view of the role of the market mechanism i.e. the communication protocols required to provide the market dialogues, the dynamics of various attributes as externally imposed market parameters vary etc. The thesis explores the above question in depth with particular focus on the double auction (or continuous double auction CDA) market. Through simulation, the oft-reported advantages of the CDA are evaluated. The thesis also briefly examines the Quote Driven market as an alternative mechanism to the CDA. The Quote Driven market has been advocated/implemented in financial markets alongside auction based systems [SEBI 2000; Theissen 2000], and a comparison between auction or order driven markets has been a topic of hot debate with economists and regulatory governing bodies. However, the same cannot be said of multi-agent systems and automated markets. This thesis contributes to the study of

7 http://www.research.ibm.com/infoecon/
8 http://www-uk.hpl.hp.com/esm/

alternative market mechanisms by offering a comparative assessment of auction protocols, which may provide new insights into the trade-offs between different types of automated markets.

Although auctions have served as the primary mode for implementing automated market mechanisms [Sandholm 1999], they have their downsides in that they might not be collusion proof (may rely on the integrity of the auctioneer); may require centralised control; and may require synchronization of many actors. In fact the optimal allocation of resources in an auction market is critically dependant on the ability of the market to aggregate demand and supply. The need for continuous trading implies that it becomes difficult to establish a critical mass of buyers and sellers. In such a scenario it is useful to study alternative mechanisms, which help to achieve continuity as well as providing liquidity (the ability to influence a trade instantaneously without affecting the stability of the market unduly). In our work [Bourne and Zaidi 2001; Zaidi, Bourne et al. 2002] we have argued that providing incentives to third party traders (market makers in a Quote Driven market) can lead to a better overall service for clients wishing to trade goods in real time. This work has established a specification of market roles, protocols and infrastructure as well as suitable evaluation metrics for a comparative assessment of automated electronic markets [Bourne and Zaidi 2001] which is expanded upon in this thesis.

The research also establishes a point, which has come into prominence in the context of emerging electronic marketplaces; namely the need for mechanisms that provide adequate profit making potential for 'third party' providers or electronic intermediaries [Foss, Garcha et al. 2000] besides the buyers and sellers. Such mechanisms will be of considerable importance in an automated environment where the need for intermediaries providing brokerage services is considerably enhanced.

1.6 Main Contributions

- An in-depth survey of automated market mechanisms, bringing together contributions from economics and multi-agent systems. This work provides a critical exposition of the underlying concepts derived from these and other disciplines and builds a comprehensive picture of the state of the art in the field of automated markets.

- Investigation of the properties of automated market mechanisms including a detailed study of the current research into the Continuous Double Auction Markets (CDA). This is supported by a simulation study, which examines the dynamics of a double auction market from various perspectives including buyers and sellers and the market as a whole (system design and performance).
- Novel results obtained from the simulation study indicate that the efficiency of the CDA market is not absolute. It is in fact dependant on a high degree of synchronization between buyers and sellers.
- The thesis explores the need for an alternative market mechanism to the CDA. In this context a Quote Driven market is presented. This work has developed the necessary market protocols and trading infrastructure in the form of a generic Quote Driven market[9]. The exposition of a Quote Driven market in the context of multi agent systems is novel and has not been attempted before.
- A Simulation infrastructure is detailed which can be utilised to undertake a comparative assessment of the Quote Driven market against the more widely used CDA market. This has not been undertaken by other researchers to-date and could provide valuable trade offs for an alternative market mechanism.
- Using the above market mechanism(s), the thesis examines case studies involving real world scenario(s) of energy and bandwidth trading.

1.7 Organization of the Thesis

This chapter provided an overview of the context for this research. The following chapters provide a more exhaustive picture of the relevant literature, applications and the state of the art in agent based trading in automated markets.

Chapter 2 provides a review of the important concepts from economic theory that can be applied to the design of agents and the protocols for interaction amongst them specifically in settings where multiple buyers/sellers are in contact (i.e. market mechanisms). It also analyses the problem of negotiation among computational entities or agents. The author examines the criteria for comparing such mechanisms and these are discussed specifically for a class of mechanisms known as auctions.

Chapter 3 is an excursion into automated trading. It builds up from the literature in experimental markets and discusses work in agent based double auction

[9] The formulation of this concept within an agent context is novel and has never been attempted previously.

markets. The Quote Driven model is also introduced in this chapter, and is examined in the light of an alternative to the double auction. The chapter discusses the organization of financial markets in general where the Quote Driven markets have found wide acceptance and application over the years. Finally it also considers example applications where agent technology can be useful. A Quote Driven market model applied to bandwidth trading in a telecommunication network is presented.

Chapter 4 discusses the experimental analysis performed and the results obtained from the simulation of a Double auction market. It also provides a description of a comparative setup for the Quote Driven market.

Chapter 5 provides a discussion of the direction of the research undertaken and possible areas for further exploration. Finally, the chapter ends with a summarized conclusion of the work undertaken.

Chapter 2
Background

This chapter is broadly structured into two parts. The *first section* presents ideas from economic theory which concern how self-interested agents (individuals) behave in the real world and their interactions in a distributed open environment or through established institutions such as markets etc. These issues are very pertinent in the design of self-interested software agents, particularly the mechanisms/protocols for their interaction, which establish globally desirable outcomes. Economic theory has a well established literature in modeling human choice and behaviour (individual *Decision Theory*), *Game Theory* (describes situations where groups of humans interact), *General Equilibrium Theory* (deals with trade and production and the interaction of a large number of consumers and producers through established institutions as markets) and *Mechanism Design Theory* (studies the design of efficient mechanisms i.e. the relationship between market structures and efficiency of outcomes). The *second section* looks at the notion of agents and Multi Agent Systems (MAS) as they are studied in Artificial Intelligence (AI). We discuss interaction and negotiation within a MAS context and the protocols available to us as system designers.

2.1 Classical Economics

2.1.1 Basics

2.1.1a Agents, Behaviour, Institutions and Equilibrium

Within economics, microeconomic theory [Kreps 1990; Mas-Colell 1995] concerns itself with the behaviour of individual economic agents and the aggregation of their actions in different institutional frameworks. The section below is an overview of these concepts.

1) *Agent:* The term agent (or actor) in economic theory refers to a decision-making entity (a human being), which has strict preferences over a consumption bundle (available commodities) within limits of its budget

constraints. The preference relation is represented mathematically by a utility function.

2) *Behaviour:* The market economy essentially captures the desire of an agent to maximize (or profit) through exchange. The decision-making agent is faced with the option of exchanging one commodity for another (at a price) so that it maximizes some objective function (the utility function).

3) The *Institution:* The actions of an individual in a market economy are governed by the opportunities available to him/her as well as the collective actions of the other interacting agents. In microeconomic theory, prices in a marketplace constitute the institutional framework.

4) *Equilibrium:* A price (or vector of prices) for which the supply meets demand for a particular commodity (or commodities) at which the aggregate excess demand is zero. The single commodity case is referred to as partial equilibrium whereas the result obtained for the entire set of commodities in the marketplace is referred to as general equilibrium.

2.1.1b Demand, Supply and the Price Process

The allocation of scarce resources is a topic that has long been studied in economics. A scarce resource can be defined as one for which the demand at a price of zero would exceed the available supply. Microeconomics is primarily concerned with the processes by which allocation of scarce resources occurs among alternative users, and of the role of prices and markets in this process.

A market can be defined as a set of arrangements by which buyers and sellers are in contact to exchange goods or services. The quantity of a commodity (goods or service) that buyers are prepared to purchase at each possible price is referred to as demand, and the quantity of a commodity that sellers are prepared to sell at each possible price is referred to as supply. In general, the greater the price of a commodity, fewer buyers will be inclined to make a purchase, and so the quantity demanded reduces. Thus a plot of price as a function of quantity has the demand curve sloping downward. In contrast the greater the price of a commodity, the more sellers are inclined to sell, and so the quantity supplied increases. Hence on a plot of price as a function of quantity, the supply curve slopes upwards. From these considerations, it is clear that at high prices the quantity supplied may exceed the quantity demanded (i.e., there is a surplus, or excess supply), and at low prices the reverse may be true (giving a shortage, or excess demand). But, at some intermediate price, the quantity demanded is equal to the quantity supplied: this is the equilibrium price, which '*clears the market*'. Graphically, the equilibrium price (and quantity) can be determined by the intersection of the supply and demand curves, as illustrated in Figure 2.1 below.

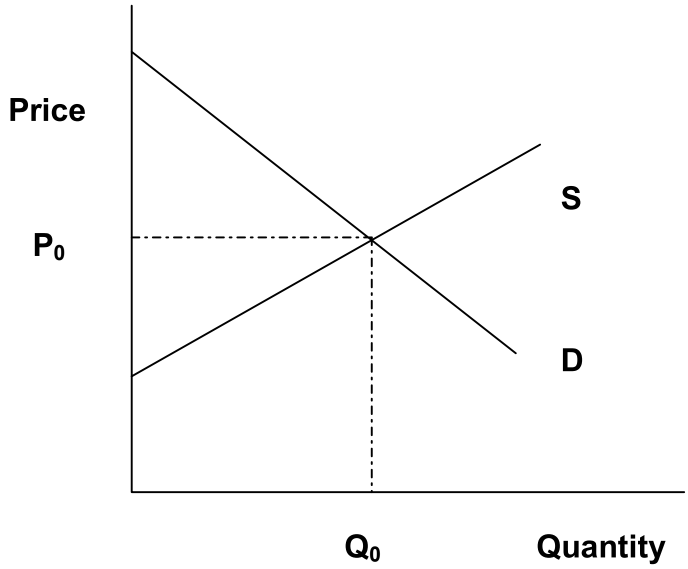

Figure 2.1: An illustration of supply and demand. The Supply Curve **S** slopes upwards and the Demand Curve **D** slopes downward. The two curves intersect at a point indicating the Equilibrium Price P_0 and the Equilibrium (Clearing) Quantity Q_0.

2.1.2 The Role of the Trading Mechanism

The supply and demand intersection explains the price of a commodity perfectly in equilibrium. However, in most standard economic theory texts, the question of how exactly this equilibrium is achieved is not considered [O'Hara 1995]. The nineteenth century economist Mary Walras suggested a probable mechanism for determining the equilibrium in a market. There are usually two interpretations of the above process in literature, which have for long been considered by economists as descriptive of the actual process to achieve equilibrium.

1) The *tatonnement* (groping) process wherein a disinterested participant announces a price (for a certain commodity) to which traders respond with quantity messages specifying their net trades at that price. The auctioneer adjusts the price repeatedly in response to the trader's messages until the

quantities equate to zero. At this point, the net trades announces by the traders are executed at the clearing price.

2) *A one-shot Clearing House mechanism*: The trader messages in this mechanism are just their demand-supply curves over possible prices. The auctioneer (disinterested participant) determines the final allocation, which arises as a result of the maximal net trade occurring at a single market-clearing price.

Although the *tatonnement* process (or the clearing house mechanism) has been used within computer science for distributed resource allocation [Cheng and Wellman 1998] and is also the process by which electricity trading is organized in pools, it has several problems. The inherent assumptions about rationality of traders, complete information about prices (i.e. effects of time and uncertainty are neglected) etc. are not plausible in open distributed marketplaces. Moreover the mechanism is susceptible to manipulation as truth telling is not a *Nash Equilibrium*. The economist Friederik Hayek captures the essence of the problem below;

”*The problem is in no way solved if we can show that all facts, if they were known to a single mind, would uniquely determine the solution; instead we must show how a solution is produced by the interaction of people each of whom possesses only partial knowledge. To assume all knowledge to be given to a single mind in which we assume it to be given to us as the explaining economists is to assume the problem away and to disregard everything that is significant in the real world*.” [Hayek, Amer. Econ. Rev. 35(4):pp 5 30,1945]

The lack of the Walrasian mechanism in explaining the market processes is evident in that very few economic institutions really follow the model in the real world. Hence there is much emphasis now on the importance of actual trading mechanisms. This is seen in the development of relatively new branches in Economics, i.e. *Market Microstructure* and *Experimental Economics* [see chapter 3].

The design of a trading mechanism is the most important determinant of market performance. Recent changes in major European stock markets and in the trading mechanisms in respect of commodity trading for electricity and bandwidth show that the question of which mechanism is the best suited is far from resolved. For example, in respect of stock exchanges, the London Stock Exchange replaced the Quote Driven trading system with the electronic order-driven system SETS in 1997. On NASDAQ, public limit orders now compete with dealer quotations. In France and Germany on the other hand, dealers were introduced to provide additional liquidity to the electronic continuous auction markets NSC and XETRA, respectively [Theissen 2000]. The changes in the electricity trading system in England & Wales as exemplified in the adoption of NETA[10] (New Electricity Trading Arrangements) and the failure of the decentralised electricity market in California [Wilson 1999] are also indicative of the necessity of more empirical research into the relative advantages of the principal trading mechanisms. The replacement of the old mechanism of electricity trading in pools by the bilateral trading model under NETA was undertaken because the Pool trading mechanism was susceptible to manipulation by the generators and led to higher wholesale prices (various OFGEM publications concerning electricity trading arrangements 1998 onwards[11]). An interesting aspect of this transformation is that NETA has replaced the '*uniform price*' auction (i.e. successful bidders all receive the same price for multiple units of output, with the price being the highest bid price accepted) in the Pool's day ahead market with a '*discriminatory*' auction (i.e. successful bidders receiving prices for each unit of output equal to the prices they actually bid) in the bilateral market. Although NETA has been in operation since Feb 2001, simulation studies comparing the two models (prior to the implementation of NETA) have suggested that in the long run the modified market arrangement may not be inherently advantageous in itself [Bower and Bunn 1999].

The next section categorizes the trading institutions, which exist in the real world and/or have been modelled and studied in laboratories. In depth discussion on the workings of these markets are left for later chapters.

2.1.3 Trading Institutions

The large variety of trading institutions seen worldwide can be classified according to several criteria. One broad classification, which distinguishes among types by the timing of decisions made by the actors involved, is given under *tables 2.1* and *2.2*. In simple environments, the decisions can be made independently (or simultaneously); this is summarized in *table 2.1.* More complex institutions feature decisions made sequentially and in real time (see *table 2.2*).

[10] To be replaced by BETTA which extends NETA to Scotland in addition to England & Wales currently.

[11] http://www.ofgem.gov.uk/public/pubframe.htm

Table 2.1 *Trading Institutions with Simultaneous Decisions (from [Davis and Holt 1993])*

	#Sellers/# Buyers (# units)	**Who Proposes Prices**	**Decisions and Timing**	**How Contracts Confirmed**
POSTED OFFER AUCTION	-/-	Sellers	Offers posted simultaneously	Buyers shop in sequence
Ultimatum Bargaining	1/1	Seller	Seller makes single offer on one unit	Buyer accepts or rejects
POSTED BID AUCTION	-/-	Buyers	Bids posted simultaneously	Sellers shop in sequence
Discriminative Auction	1/- (N units)	Buyers	Bids posted simultaneously	Highest N bidders pay own bids
1st Price Sealed-Bid Auction	1/- (1 unit)	Buyers	Bids posted simultaneously	High bidder pays own "1st" price
Competitive Sealed-Bid Auction	1/- (N units)	Buyers	Bids posted simultaneously	Highest N bidders pay N+1st price
Second Price Sealed-Bid Auction	1/- (1 unit)	Buyers	Bids posted simultaneously	Highest bidder pays 2nd price
Clearinghouse Auction	-/-	Buyers and sellers	Bids posted simultaneously	Intersection of bid and offer arrays

Table 2.2 *Trading Institutions with Sequential Decisions (from [Davis and Holt 1993])*

	#Sellers/# Buyers (# units)	**Who Proposes Prices**	**Decisions and Timing**	**How Contracts Confirmed**
Dutch Auction	1/- (1 unit)	Seller clock	Price lowered sequentially	Buyer who stops clock
English Auction	1/- (1 unit)	Auctioneer	Price raised sequentially	Sale to high bidder
Bid Auction	-/-	Buyers	Price raised sequentially	Sellers
Offer Auction	-/-	Sellers	Price lowered sequentially	Buyers
Double Auction	-/-	Both types	Bids raised and offers lowered sequentially	Both types
Decentralized Negotiation	-/-	Both types	Sequential but decentralized	Both types

Another classification of trading institutions is more specific to trading in financial markets. This classification [Madhavan 2000] is undertaken on the basis of:

1) *Degree of Continuity*
   a) Periodic or Batch Markets: Allow trading at only specific points in time.
   b) Continuous Markets: Allow trading at any point in time (when market is open).
2) *Reliance on Market Makers*
   a) Auction (or order driven) markets: Trades occur between public investors without any dealer intermediation e.g. through a public order book.
   b) Dealer (or Quote Driven) markets: A market maker (or dealer) takes the opposite side of every transaction.

Figure 2.3 *Variations in Real-World Financial Markets*

	NASDAQ NMS	NYSE Open	NYSE Intraday	Paris Bourse	POSIT	CBOT	FX Market
Market Type							
Continuous	X		X	X		X	X
Dealer Presence	X	X	X			X	
Multilateral		X			X		
Transparency							
Pre-Trade Quotes	X		X	X		X	
Post-Trade Quotes	X	X	X	X	X	X	

2.1.4 Game Theory

Classically, in equilibrium markets, the agents are assumed to act *competitively* i.e. they treat prices as exogenous. Each agent reveals its demand (supply) decisions truthfully so as to maximise its utility (profits) given the prevailing market prices, assuming that it has no impact on those prices. This assumption holds if the market is so large that no single agent's actions affect the prices i.e. the number of agents approaches infinity. However if the number of agents in a market is finite, an agent can act strategically and potentially achieve higher utility by over/under representing. The branch of economics, which studies the interaction of decision makers who are conscious that their actions affect each other is known as Game Theory [Fudenberg and Tirole 1991; Mas-Colell 1995; Rasmusen 2001]. Game Theory is essentially a mathematical model of interaction of rational agents and deals with issues like conflict, coordination and/or cooperation.

2.1.4a Basic Definitions

The essential elements of a game are **players**, **actions**, **payoffs** and **information**. These are collectively known as the **rules of the game**. In trying to maximize their payoffs, the players will devise plans known as **strategies** that pick actions depending on the information that has arrived at each moment. The combination of strategies chosen by each player is known as the **equilibrium**. The actions coming out of all the player's plans, gives the **outcome** of a game.

The utility of an agent is described (as in expected utility theory) as its preferences over a set *O* of outcomes. The fundamental concept of agent choice in a game is known as a strategy.

Definition: Player *i's* **strategy** s_i is a rule that tells him which action to choose at each instant of the game given his information set.

Player *i's* **strategy set** or **strategy space** $S_i = \{ s_i \}$ is the set of strategies available to him. A **strategy combination** s = $(s_1,\ldots,s_n)$ is an ordered set consisting of one strategy for each of the n players in a game.

Definition: Player *i's* **payoff** $\Pi_i (s_1,\ldots,s_n)$ means either:

1) The utility player i receives after all players and nature (pseudo player) have picked their strategies and the game has been played out; or
2) The expected utility he receives as a function of the strategies chosen by him and the other players.

Definition: An **equilibrium** $s^*= (s_1^*,\ldots,s_n^*)$ is a strategy combination consisting of a best strategy for each of the n players in the game.

An equilibrium or solution concept defines an equilibrium based on the possible strategy combinations and the payoff functions.

2.1.4b Solution Concepts

Game theory provides a number of solution concepts to compute the outcome of a game with self-interested agents, given assumptions about agent preferences, rationality, and information available to agents about each other.

It is useful to introduce notation $s_{-i} = (s_1,\ldots,s_{i-1}, s_{i+1},\ldots s_n)$ for the strategy of every agent except agent i.

Definition: Player *i's* **best response** or **best reply** to the strategies s_{-i} chosen by the other players is the strategy s_i^* that yields him the greatest payoff; i.e.

$$\Pi_i (s_i^*, s_{-i}) \geq \Pi_i (s'_i, s_{-i}) \quad \text{for all } s'_i \neq s_i^*$$

The first important solution concept is dominant strategy equilibrium.

Definition: **[dominant strategy]** The strategy s_i^* is a **dominant strategy** if it is a player's strictly best response to any strategies the other players might pick, in the sense that whatever strategies they pick, his payoff is the highest with s_i*. Mathematically:

$$\Pi_i (s_i*, s_{-i}) > \Pi_i (s'_i, s_{-i}) \quad \text{for all } s_{-i}\text{, for all } s'_i \neq s_i*$$

A **dominant strategy equilibrium** is a strategy combination consisting of each player's dominant strategy.

The representation below is from the widely studied game, *The Prisoner's Dilemma*. It helps explain the notion of dominant strategy.

B

A

	Deny	*Confess*
Deny	-1, -1	-10, 0
Confess	0, -10	-8, -8

Payoffs to: (A, B)

Consider two prisoners say **A** and **B** who are being interrogated by police separately. They have two options, *Deny* or *Confess* to their crime. The payoffs (i.e. the number of years of imprisonment) for all cases are given in the corresponding row and column. In the above setting, each player has a dominant strategy. Consider prisoner A (he does not know what B is choosing), if B chooses *Deny*, A faces a *Deny* payoff of –1 and a *Confess* payoff of 0, whereas if B chooses Confess, A faces a Deny payoff of –10 and a Confess payoff of –8. In either case A does better with *Confess*. The similar argument holds for B. Therefore the dominant strategy equilibrium is (*Confess, Confess*) and the equilibrium payoffs are (-8, -8). Note that this is worse for both players than (-1, -1) which could be achieved if both *Deny*. The two could possibly be better off with *Deny*, but that is only possible if the two prisoners could talk to each other and make binding commitments to each other with respect to the decisions they would take. Within game theory, such games where binding commitments are possible are known as **cooperative games** whereas our main concern (as the treatment of the game above) is to do with **non-cooperative game theory** where binding commitments are not possible.

Dominant strategy equilibrium provides a very robust solution concept independent of the information structure of the game and is very useful for implementation of mechanism design considerations. However there are very few instances where dominant strategy equilibrium is actually present in a game. Generally, in such cases the concept of Nash equilibrium is most widely used.

Definition: **[Nash equilibrium]** A strategy combination s^* is a **Nash equilibrium** if no player has incentive to deviate from his strategy given that other players do not deviate. Formally,

For each i $\quad \Pi_i (s_i^*, s^*_{-i}) \geq \Pi_i (s'_i, s^*_{-i})$ for all s'_i

In words, every agent maximizes its utility with strategy s_i, given its preferences and the strategy of every other agent. The definition of Nash equilibrium lacks the "for all s_{-i} " of dominant strategy equilibrium, so a Nash strategy need only be the best response to other Nash strategies and not to all possible strategies.

The above solution concepts explicitly assume that the players know all relevant information about each other (including the payoffs that each receives from the various outcomes). Such games are known as games of *complete information*. However, such explicit knowledge about other player's utilities or payoffs is hardly available in the real world. Rather in many circumstances players have *incomplete information*. In such cases, the approach is to consider each player's preferences as being determined by the realization of a random variable. Although the random variable's actual realization is observed only by the player, its ex ante probability distribution is assumed to be common knowledge among all the players. This formulation allows reinterpretation of a game of incomplete information as a game of *imperfect information*. Nature is assumed to make the first move in such games by choosing realizations of the random variables that determine each player's preference **type,** θ (the strategy set, information partition and payoff function). Such formulations are referred to as **Bayesian** games and the solution concept proposed for such games is referred to as Bayesian (or **Bayesian-Nash**) equilibrium. A Bayesian equilibrium is a Nash equilibrium where every agent is assumed to share common prior (prior beliefs) about the distribution of agent types, F(θ), such that for any particular game the agent profiles are distributed according to F(θ). These prior beliefs are updated during the course of the game depending on the observed actions of the other players (which are assumed to be following equilibrium behaviour). The updating is done according to Baye's rule, which represents a standard way to handle imperfect information.

Comparing Bayesian-Nash with Nash equilibrium, the key difference is that agent *i's* strategy $s_i(\theta)$ must be a best-response to the distribution of strategies of other agents, given distributional information about the preferences of other agents. Agent *i* does not necessarily play a best-response to the actual strategies of the other agents.

Bayesian-Nash makes more reasonable assumptions about agent information than Nash, but is a weaker solution concept than dominant strategy equilibrium. Remaining problems with Bayesian-Nash include the existence of multiple equilibria, information asymmetries, and rationality assumptions, including common-knowledge of rationality.

The solution concepts of Nash, dominant-strategy and Bayesian-Nash hold in both **static** and **dynamic** games. In a static game every agent commits to its strategy simultaneously (think of a sealed-bid auction for a simple example). In a dynamic game, actions are interleaved with observation and agents can learn information about the preferences of other agents during the course of the game (think of an iterative auction, or stages in a negotiation). Additional refinements to these solution concepts have been proposed to solve dynamic games, for example to enforce sequential rationality (backwards induction) and to remove non-credible threats off the equilibrium path. One such refinement is sub-game perfect Nash, another is perfect Bayesian-Nash (which applies to dynamic games of incomplete information), see [Fudenberg and Tirole 1991] for more details.

The next section is an introduction to mechanism design theory. The game theoretic solution concepts discussed above provide the basis for the discussion that follows. An ideal mechanism provides agents with a dominant strategy and also implements a solution to the multi-agent distributed optimization problem. We can state the following preference ordering across implementation concepts: Dominant > Bayesian-Nash > Nash.

2.1.5 Mechanism Design: Important Concepts

The mechanism design problem is to implement an optimal system-wide solution to a decentralized optimisation problem with self-interested agents with private information about their preferences for different outcomes

The system-wide goal in mechanism design is designed with a *social choice function*, which selects the optimal outcome given agent types.

Definition: **[Social choice function]** A Social choice function

$$f : \theta_1 \times \theta_2 \times \ldots\ldots \times \theta_I \rightarrow O$$

chooses an outcome $f(\theta) \in O$, given agent types $\theta = (\theta_1,\theta_2, \ldots\ldots\ldots \theta_I)$.

In other words, given agent types $\theta = (\theta_1,\theta_2, \ldots\ldots\ldots,\theta_I)$ we would like to choose outcome $f(\theta)$. The mechanism design problem is to implement "rules of a game", for example defining possible strategies and the method used to select an outcome based on agent strategies, to implement the solution to the social choice function despite an agent's self-interest.

Definition: **[mechanism]** A mechanism $M = (\sum_1,\ldots\ldots,\sum_I, g(\bullet))$ defines the set of strategies $\sum_i$ available to each agent, and an *outcome rule g*: $(\sum_1 x \ldots x \sum_I) \rightarrow O$, such that g(s) is the outcome implemented by the mechanism for strategy profile $s = (s_1,\ldots..,s_I)$.

In words, a mechanism defines the strategies available (e.g., bid at least the ask price, etc.) and the method used to select the final outcome based on agent strategies (e.g., the price increases until only one agent bids, then the item is sold to that agent for its bid price).

Game theory is used to analyze the outcome of a mechanism. Given mechanism M with outcome function $g(\bullet)$, we say that a mechanism implements social choice function $f(\theta)$ if the outcome computed with equilibrium agent strategies is a solution to the social choice function for all possible agent preferences.

Definition: **[mechanism implementation]** Mechanism $M=(\sum_1,\ldots\ldots,\sum_I, g(\bullet))$ implements social choice function

$f(\theta)$ if $g(s^*_1(\theta_1),\ldots\ldots, s^*_I(\theta_I)) = f(\theta)$ for all $(\theta_1,\ldots,\theta_I) \in \theta_I \times \theta_2 \times \ldots.. x\ \theta_I$ where strategy profile $(s^*_1,\ldots.., s^*_I)$ is an equilibrium solution to the game induced by M.

The equilibrium concept may be Nash, Bayesian-Nash, dominant or some other concept; generally as strong an equilibrium solution as possible.

To understand the difficulty with the mechanism design problem, consider a very naïve mechanism, and suppose that the system-wide goal is to implement social choice function $f(\theta)$. The mechanism asks agents to report their types, and then simply implements the solution to the social choice function that corresponds with their reports, i.e. the outcome rule is equivalent to the social choice function, $g(\theta) = f(\theta)$ given reported types $\theta = (\theta_1,\theta_2, \ldots\ldots\ldots \theta_I)$. But, there is no reason for agents to report their true types. In a Bayesian-Nash equilibrium each agent will choose to announce a type θ_i' to maximize its expected utility given distributional information about the types of other agents, and under the assumption that the other agents are also following expected-utility maximizing strategies. This announced type θ_i' need not equal the agent's true type.

Looking ahead, the mechanism design problem is to design a mechanism - a set of possible agent strategies and an outcome rule - to implement a social choice function with desirable properties, in as strong a solution concept as possible; i.e. dominant is preferred to Bayesian-Nash because it makes fewer assumptions about agents.

Many properties of a mechanism are stated in terms of the properties of the social choice function that the mechanism implements. A few of the desirable properties for social choice functions are:

1) Pareto optimality (or Pareto efficiency)
2) Allocative efficiency (or efficiency)
3) Budget-Balance property

A discussion of these properties is deferred for the next chapter, which discusses negotiation mechanisms in agents, in particular auction protocols.

2.2 Negotiation in Multi Agent Systems

The lifecycle of a typical business process (based on [Jennings, Faratin et al. 1996]) is constructed below:

1) *Matchmaking*: A trader locates other traders that it could potentially do business with.
2) *Negotiation*: The trader enters into negotiation with one or more of these potential business partners, to see if they can agree mutually acceptable terms of business. These terms could include a definition of the goods or service being traded, price, delivery date, etc.
3) *Contract Formation*: These agreed terms are placed into a legally binding contract.
4) *Contract Fulfilment*: The parties carry out the agreed transaction, within the parameters specified in the contract.

Agent technology has been applied with considerable success to all the above processes, but for the purpose of this thesis, (as also generally the field of automated markets) we concern ourselves primarily with the Matchmaking and Negotiation phases. In auction protocols, the matchmaking is carried out by a third (disinterested) party known as the auctioneer. An auctioneer acts as a market intermediary providing a service to both sides (buyers and sellers) of a market. In general the primary roles of a market intermediary [Bailey and Bakos 1997] can be:

1) To *aggregate* buyer demand and seller products to achieve economies of scale and to reduce bargaining asymmetry.
2) Protect buyers and sellers from the opportunistic behaviour of other participants in a market by becoming an agent of *trust*.
3) *Facilitate* the market by reducing operating costs and,
4) *Match* buyers and sellers.

This section introduces the notion of agents and multi agent systems (MAS) in particular and expands on the needs and requirements for automated negotiation in such a context. This is followed by a description of the criteria for evaluating different negotiation protocols in multi agent systems.

2.2.1 Agents and Multi-Agent Systems

The notion of an '*agent*' has developed independently in several disciplines notably AI, Economics and as a software paradigm. This has meant that though the agent community has benefited from interdisciplinary research in the field, there is no universally acceptable definition of the term itself. Researchers tend to associate numerous attributes to the term 'agent' including autonomous behaviour, intelligence, mobility, learning/adaptability etc. These attributes will have varying importance in differing domains and some may even have additional characteristics. However for most practical purposes (there is broad agreement on autonomous behaviour being a qualifying characteristic of agent behaviour) the definition below will serve to illustrate the notion of agency.

Definition: [Agent] A computer system situated in some environment and capable of autonomous action in this environment in order to meet its design objectives.

Much of traditional AI has been concerned with the development of individual intelligent entities or agents, with a single locus of internal reasoning and control implemented in a Von Neumann architecture. But intelligent agents do not function in isolation – they are at the least part of an environment in which they operate and the environment typically contains other such intelligent systems. Thus it makes sense to view such systems in societal terms. This is the area, which Distributed Artificial Intelligence (DAI) looks at.

Definition: [Distributed Artificial Intelligence (DAI)] is the study, construction and application of multi agent systems, that is, systems in which several interacting, intelligent agents pursue some set of goals or perform certain tasks. The research within DAI can be classified into two domains:

1. **Cooperative Distributed Problem Solving (CDPS)**: The system designer imposes an interaction protocol and a strategy (a mapping from state history to action; a way to use the protocol) for each agent. The main question is what social outcomes follow given the protocol and assuming that the agents use the imposed strategies.
2. **Multi Agent Systems (MAS)**: In MAS, the agents are provided with an interaction protocol, but each agent will choose its own strategy. A self-interested agent will choose the strategy that is best for itself, which cannot be explicitly imposed from outside. Therefore the protocols need to be designed with a non-cooperative strategic perspective: the main question is what social outcomes follow given a protocol which guarantees that each agent's desired local strategy is best for that agent – and thus the agent will use it.

The characterization of agents as selfishly and autonomously pursuing multiple goals has a number of important implications. The pursuit of individual goals is beneficial in that it decouples agents from one another. Thus, self-interest, as a behaviour guideline, encourages separation between individual and group problem solving. Also, the assumption in MAS that agents may have multiple, and at least partially, conflicting goals produces social dilemmas or real conflict. In such circumstances there is a need for the agents to interact in a meaningful manner (i.e. coordinate or negotiate amongst themselves) in order to achieve their goals. Therefore the approach outlined above for MAS is required for the design of robust non-manipulable multi agent systems in open distributed environments where the agents may represent different real world parties. This research is concerned with addressing some of the issues that arise within the context of multi agent systems. We take Electronic Commerce as an exemplar of a system, which incorporates interaction between computational components. In particular we address ourselves with the issue of trading interactions among computational agents that represent buyers and sellers in open marketplaces.

2.2.2 Rationale for Co-ordination

The need for coordinating interactions among agents arises because of several reasons. These can be classified into two broad categories [Faratin 2000]:

1) *Coordination as a function to inform local activities*: This is seen in distributed problem solving applications when there are interdependencies between agents' actions, between local actions and some global criteria that need to be satisfied, or when there are differences in expertise or levels of resources.
2) *Coordination to solve conflicts of interest*: This is seen in the field of multi agent systems where the individual agents may have mutually exclusive goals (i.e. a buyer agent wants to buy a commodity cheaply whereas a seller agent would like to sell at the highest possible price). Coordination in such circumstances is motivated by the desire to make a *deal* while selfishly maximising personal goals. Such coordination is referred to as negotiation.

Definition: [Negotiation] A process by which a joint decision is made by two or more parties. The parties first verbalize contradictory demands and then move towards agreement by a process of concession making or search for new alternatives.

2.2.3 Evaluation Criteria for Negotiation Mechanisms

Negotiation mechanisms can be evaluated according to many types of criteria [Sandholm 1999].

1) *Social Welfare*: This represents the sum of all the agent's payoffs or utilities in a given solution, i.e. it serves as a measure of the global good of the agents. It can be used as a criterion for comparing alternative mechanisms by comparing the solutions that the mechanisms lead to.
2) *Pareto Efficiency*: A solution x is Pareto efficient (or Pareto optimal) if there is no other solution x' such that at least one agent is better off in x' than in x and no agent is worse off in x' than in x. Pareto efficiency also measures the global good.

3) *Individual Rationality*: Participation in a negotiation is individually rational to an agent if the agent's payoff in the negotiated solution is no less than the payoff that the agent would get by not participating in the negotiation. A mechanism is individually rational if participation is individually rational for all the agents.

4) *Stability*: A mechanism designed for self-interested agents should be stable (non-manipulable) i.e. it should motivate the agent to behave in a desired manner. Such mechanisms could have dominant strategies for the agents. More commonly Nash equilibrium or some refinement is used.

5) *Computational Efficiency*: Mechanisms should be designed so that the agents utilising them have to expend the least amount of computation as possible.

6) *Distribution and Communication Efficiency*: Among mechanisms, all the above properties being equal, preference is given to distributed protocols which avoid singular points of failure.

Another important property of mechanisms is incentive compatibility, which is discussed below for the class of negotiation mechanisms known as direct revelation mechanisms.

2.2.3a The Revelation Principle, Incentive-Compatibility, and Direct-Revelation

The revelation principle states that under quite weak conditions any mechanism can be transformed into an equivalent incentive-compatible direct-revelation mechanism, such that it implements the same social-choice function. This proves to be a powerful theoretic tool, leading to the central possibility and impossibility results of mechanism design.

A direct-revelation mechanism is a mechanism in which the only actions available to agents are to make direct claims about their preferences to the mechanism. An incentive compatible mechanism is a direct-revelation mechanism in which agents report truthful information about their preferences in equilibrium. Incentive-compatibility captures the essence of designing a mechanism to overcome the self-interest of agents. In an incentive-compatible mechanism an agent will choose to report its private information truthfully, out of its own self-interest. As an example, the second-price sealed-bid (or Vickrey) auction, discussed later, is an incentive-compatible (actually strategy-proof i.e. agents do not need to model the preferences and strategies of other agents) direct-revelation mechanism for the single-item allocation problem.

2.2.4 Coordination Mechanisms in MAS

There are several coordination (negotiation) mechanisms studied in multi agent systems. These include voting, auctions, bargaining, markets, contracting and coalition forming [Sandholm 1999]. For the purposes of relevance of this research, we will concentrate here on auctions mechanisms.

2.2.4a Auction Theory

Auction theory analyses protocols and agent strategies in auctions. As indicated previously (section 1.3), auctions provide several advantages and have become very important in the context of trading in e-commerce systems of the day. Auctions have also been used for long within the computer science fraternity in problems ranging from distributed scheduling [Wellman, Walsh et al. 2001], CDPS and more recently within multi agent systems.

2.2.4b Auction Settings

Depending on how an agent's value of an item is formed in an auction, we can define three different settings for the auction.

1) *Private value auction*: The value of the commodity depends only on the agent's own preferences i.e. auctioning a cake that the winning bidder will eat. The key element is that the winner will not resell the item or get utility from showing it off to others. The agent is often assumed to know the value for the commodity exactly and this information is private to the agent.
2) *Common value auction*: The actual value of an item is the same for everyone, but bidders may have different private information about what that value actually is; e.g. auctioning treasury bills or the lease of oil-exploration rights. Nobody inherently prefers having the bills; the value of the bill comes entirely from the possibility of resale.
3) *Correlated value auction*: The agent's value depends upon its own preferences and partly on other's values. For example, bidding for a painting whose value may be a function of how much the agent likes it but also

depends on other's private information (i.e. how much they like it) because this affects resale possibilities.

Most cases of auctioning goods would fall within the last category. However the two extreme cases allow a more rigorous treatment of the subject, which is often difficult with the third case.

2.2.4c Auction Classification and Private Value Strategies

Wurman et al [Wurman 2001] define a parameterisation of the auction space by formalising a set of rules that describe an auctions policies around 3 basic tasks.

1) *Receiving bids i.e. the bidding rules*: These determine what actions participants can take.
2) *Revealing intermediate information*: Most auctions reveal information regarding their current state as the auction progresses (except sealed bid auctions) i.e. the new highest price in an English auction etc.
3) *Clearing the auction*: The task of determining prices, quantities and trading partners as a function of the bids.

The above classification can be used to define a host of auction permutations. However in this section we will be concerned with the standard auction types, which involve a single seller and multiple buyers and a single unit of the commodity. (Procurement auctions have a single buyer and multiple sellers whereas stock exchanges function as double auctions i.e. multiple buyers and sellers. The standard auction types in literature are:

1) English (or first-price open-cry) auction:

Rules: Each bidder is free to revise his bid upwards. When no bidder wishes to revise his bid further, the highest bidder wins the object and pays the bid.

Strategies: A player's strategy is his series of bids as a function of his value, his prior estimate of other player's valuation and the past bids of all the players. An agent's dominant strategy in a private value English auction is to always bid a small amount more than the current highest bid and stop when his private value is reached.

Payoffs: Winner's payoff is his value minus the highest bid. Loser's payoff is zero.

2) First Price Sealed Bid Auction:

Rules: Each bidder submits one bid, in ignorance of the other bids. The highest bidder pays his bid and wins the object.

Strategies: A player's strategy is his bid as a function of his value and his prior beliefs about other player's valuations. In general there is no dominant strategy for this auction. An agent's best strategy is to bid less than his true valuation; how much less would depend upon other's bids.

Payoffs: The winner's payoff is his value minus the bid. Loser's payoff is zero.

3) **Second Price Sealed Bid Auction (Vickrey Auction)**:

Rules: Each bidder submits one bid, in ignorance of other's bids. The bids are opened and the highest bidder pays the amount of the second highest bid and wins the object.

Strategies: A player's strategy is his bid as a function of his value and his prior belief about the other player's valuations. The dominant strategy for a player in a Vickrey auction is to bid his true value.

Payoffs: The winner's payoff is his value minus the second highest bid. Loser's payoffs are zero.

The Vickrey auction has been widely adopted for use among multi agent systems. Uses include allocation of resources in operating systems, bandwidth etc. However Vickrey auctions have several drawbacks too [Sandholm 1999] and it is sometimes advisable to choose the English auction in comparison.

4) **Dutch (or Descending) Auction**:

Rules: The seller announces a bid, which he continuously lowers until some buyer stops him and takes the object at that price.

Strategies: A player's strategy is when to stop bidding as a function of his valuation and his prior beliefs as to the other player's valuations. The Dutch auction is strategically equivalent to the first-price sealed bid auction because in both cases, an agent's bid matters only if it is the highest and no relevant information is revealed during the process.

Payoffs: The winner's payoff is his value minus the bid. Loser's payoffs are zero.

2.2.4d Efficiency of the Resulting Allocation

Given the wide variety of auctions, it is sometimes difficult to choose one over another. A lot depends upon whose purpose the auctioning mechanism serves. The seller would like to gain maximum surplus from the buyers whereas the buyers would like to acquire the item at the lowest possible price. From a mechanism design perspective, the auction should lead to socially efficient allocations; it should also reach equilibrium i.e. a state from which no participant wishes to deviate. Moreover it should be individually rational and *incentive compatible*. Incentive compatibility ensures that the agent does not spend time modelling other agents.

A striking result from auction theory is the **revenue equivalence theorem:** It states that in all the four kinds of auctions referred to above, for private independent values, the seller's expected price is the same. Also all the auction forms allocate the auctioned item pareto efficiently to the bidder who values it most.

Another result from auction theory is known as the *winner's curse*. In a common value auction setting, the object is sold to the highest bidder who is also the most likely to have overestimated the value of the object thereby ending up paying more than it's worth.

Yet another significant result essentially states that 'there is no perfect mechanism' i.e. no auction is incentive compatible, individually rational, efficient and budget-balanced at the same time [Wurman 2001].

2.2.4e Problems with auction protocols

One problem that arises with all the standard auction types is that they are not collusion proof. The bidders can coordinate their strategies to ensure that bids stay artificially low. From this perspective, the first-price sealed bid and Dutch auctions are preferable. The English and Vickrey auctions self-enforce some of the most likely collusion agreements. They are not suitable in situations where collusion can be expected.

The Vickrey auction has problems with insincere auctioneers. The auctioneer may overstate the second highest bid to the eventual winner (unless the highest bidder can verify otherwise), which would imply, that the highest bidder ends up paying more than if the auctioneer was truthful.

In non-private value auctions with the English auction protocol, besides the problem of the winner's curse, the auctioneer can use *shills* that bid in the auction in order to make the real bidders increase their valuation of the item.

The discussion of the auction mechanisms and their properties referred to above, by no means offer foolproof guidelines as to what particular mechanism to use. It does however give some guidance for mechanism designers when implementing these in specific domains and settings. One thing, which is quite evident, is the argument for the development of intelligent agent technology to represent humans in such an environment given the dynamic nature of auction interactions.

2.3 Summary

This chapter has presented the relevant concepts from economics, which are particularly significant, both in the design of software agents for participation in automated markets and in the design of the market protocols themselves. In multi agent systems there are two strands to the research into design of automated interacting systems; the *agent centric* approach studies the strategic behaviour of individual agents (how best to behave in response to a given environment) whereas the *environment centric* approach looks at the creation of agent societies in which individuals have an incentive to behave appropriately. Game theory and mechanism design provide the necessary tools to study these two approaches. It needs to be said though that the two approaches are complementary to each other. The design of *incentive compatible* mechanisms takes into account the game theoretic considerations in any interaction. Designing mechanisms to achieve specific economic requirements, such as achieving market efficiency or maximising social welfare, against self-interested intelligent traders, is no trivial matter as can be seen from accounts of the auction design process for the recent radio spectrum auctions in the UK [Klemperer 2002].

This chapter also discussed the rationale for coordination (negotiation) in multi agent systems. We identified the evaluation criteria for the broad class of such mechanisms and then addressed them for one such class i.e. auction protocols. Auctions are the most widely used mechanism for negotiation in multi agent systems. A classification of auction types is given and we have analysed their properties in implementing one-to-many negotiation situations.

Chapter 3

Automated Trading

This chapter introduces the concept of automated trading, i.e. electronic institutions where autonomous software agents following defined protocols and representing individuals/businesses interact to trade (i.e. buy and sell) goods in a distributed environment e.g. the Internet. The description summarises research into the precursor of automated markets, i.e. the experimental commodity markets studied in economics literature. This is followed by a review of the research into agent based automated markets itself. We describe in detail the two most important market organisations catering to multiple buyer/seller populations, i.e. the Double auction and the Quote Driven market. Lastly, there is a description of two domains, Electricity and Bandwidth markets which offer the kind of fine granularity where agent technology may be particularly useful. A proposed bandwidth market based on the Quote Driven model is also presented.

3.1 Experimental Economics

Historically economics has been an observational science, which seeks to explain market phenomenon generated by economic observations over time. Economists have developed sophisticated models to explain their theories but for a long time, the capacity to evaluate their predictive content was not developed. Vernon Smith[12] considers much of economic theory to be *'ecclesiastical theory'* which is accepted or rejected on the basis of authority, tradition or opinion about assumptions rather than on the basis of having survived a rigorous falsification process that can be replicated. Experimental economics (Davis and Holt, 1993, Kagel and Roth, 1993) developed as a separate field in economics in order to answer the above limitations. It attempts to create laboratory models of *microeconomic systems* [Smith 1982], which allow control of variables and replication of the processes for a better understanding and demonstrable knowledge of the economist's attempts to understand markets.

The literature in experimental economics has roughly evolved in 3 directions:

1) Market experiments

[12] www.econlabs.arizona.edu

2) Game experiments

3) Individual decision making experiments

Here we will limit ourselves to simple market experiments. The section below outlines the design of a simple market experiment, which will help to clarify some of the concepts involved. The market considered is a simple commodity market where buyers and sellers have private values for the goods.

3.1.1 A Simple Market Experiment

The table below gives an induced demand-supply schedule in a market by assigning values to buyers and costs to the sellers. There are six buyers and six sellers in the market who are allowed to trade 2 units each. Restricting the buyers to buying the higher valued units and the sellers to selling the lower cost units first enforces the downward sloping demand and the upward sloping supply characteristics as indicated in Fig. 2.1.

Buyer's Values **Seller's Costs**

Buyer	Unit 1	Unit 2	---------	Seller	Unit 1	Unit 2
B1	1.4	1.4		**S1**	1.3	1.4
B2	1.5	1.3		**S2**	1.2	1.5
B3	1.6	1.2		**S3**	1.1	1.6
B4	1.7	1.1		**S4**	1.0	1.7
B5	1.8	1.0		**S5**	0.9	1.8
B6	1.9	0.9		**S6**	0.8	1.3

Table 3.1 *Parameters for a Simple Market Experiment*

Fig 3.1 *Supply and Demand Structure for a Market Experiment*

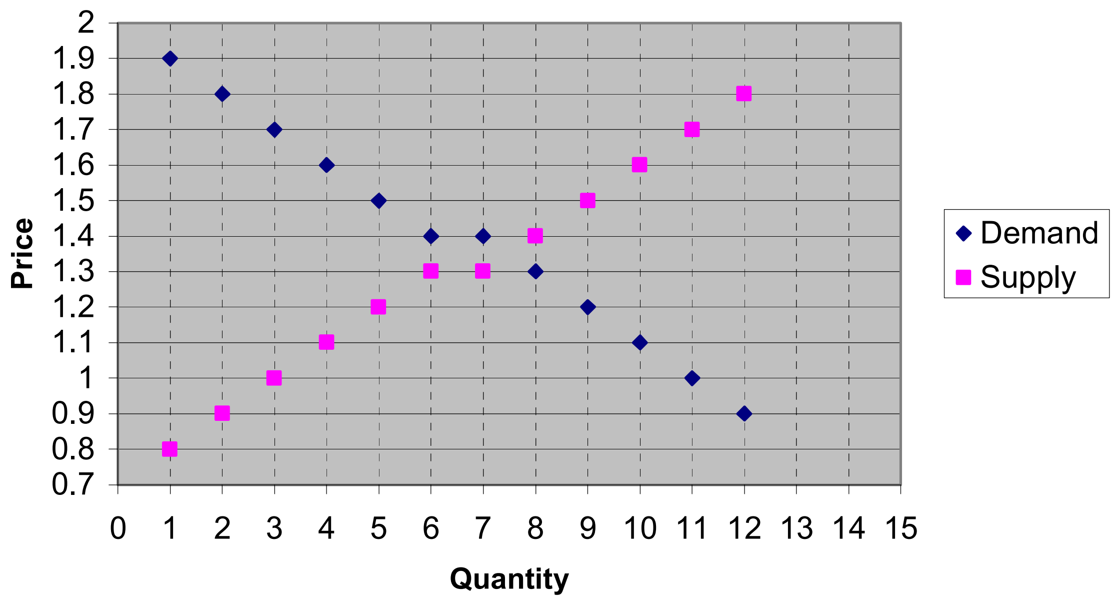

The following observables will be considered in detail for the experimentation to follow in the next chapter:

Equilibrium Price: Price at which the demand and supply curves intersect. In the figure above this price is predicted to lie in the range 1.3 to 1.4.

Equilibrium or 'Clearing Quantity': The corresponding value on the X-axis at the intersection of the 2 curves gives the predicted *competitive quantity* to be 7.

Trade Surplus: The maximum possible surplus (this gives a measure of market performance) that can be obtained by trading with the demand and supply conditions above is given as the area between the curves and to the left of the intersection (i.e. 36 in the above case). *Efficiency* (or allocative efficiency) of the market is measured as a percentage of maximum possible surplus extracted. Competitive price theory predicts 100% market efficiency with the available surplus distributed equally among buyers and sellers (i.e. 18 units to both Buyers and Sellers).

3.1.2 Market Organisation

As discussed previously, the competitive behaviour outlined above is established as a result of trading in a marketplace. The market organisation governs the information and the opportunity sets available to the agents in the market leading to a market price being established. Most markets are organised as a particular style of *auction* mechanism. Colloquially, an auction refers to set of arrangements where sellers of a commodity and potential buyers interact to agree to a price. This section discusses the auction arrangements, which allow multiple buyers and sellers to interact. These are in contrast to the one sided nature of standard auctions discussed in Chapter 2 in which a single seller receives bids from multiple buyers or multiple sellers compete for the sale of a fixed number of units or 'contracts' sought by a single buyer. Such institutions, which allow matching of demand and supply in a market with multiple buyers and sellers are employed in numerous financial institutions worldwide. The two mechanisms considered in this context are:

1) **The Double Auction**: The Double Auction (or the Continuous Double Auction **CDA**) [Friedman and Rust 1992; Davis and Holt 1993] is perhaps the most widely studied auction mechanism. This is because the markets organised under double auction trading rules appear to generate competitive outcomes much more quickly and reliably than markets organised under any alternative set of trading rules. In a Double Auction, buyers and sellers simultaneously and asynchronously announce bids and offers: at any time, a seller is free to accept the bid of a buyer and a buyer is free to accept the offer of a seller. The DA or CDA therefore resembles a parallel integration of English and Dutch auctions styles. Some DAs apply an *improvement* rule (also known as the NYSE rule), introduced to speed up the auction process that requires that each bid or offer leading to a transaction must be an improvement on the previous one.

2) **The Call Auction** (also known as the Clearing-House mechanism; see section 2.1.2): A call auction is essentially a periodic version of the Continuous Double Auction. As defined previously, a central auctioneer collects bids and offers from all buyers and sellers. This array of bids and offers is used to determine the market demand and supply curves with the intersection giving the market-clearing (or equilibrium) price. All possible trades clear simultaneously at the same price. Call auctions are employed in financial institutions worldwide especially in conditions of low liquidity or volume. This type of auction arrangement has also found wide application in trading in spot markets for electricity (i.e. the electricity pools).

The above two market structures are collectively referred to as *order-driven* markets. This is because the trading is purely driven by the arrival of customer orders, which cross each other. The other type of continuous market mechanism that facilitates trading between multiple buyers and sellers is the *Quote Driven* market, which is discussed later in the chapter.

3.1.3 Experimental Double Auction Market

Vernon Smith[13] [Smith 1962] conducted extensive studies on the double auction mechanism. In his laboratory studies, a group of human subjects are divided into sub groups of sellers (each with an entitlement to sell one or more units of a commodity at a price no lower than their specified limit price) and buyers (each with an entitlement to buy one or more units at a price no greater than their specified limit price). Each trader's individual limit price is private i.e. not known to any other trader. Each buyer is encouraged to trade in the market by being instructed to consider the difference between the given limit price and the actual contract price paid for the commodity as pure profit. Furthermore buyers are told that it is better to make no profit and own the commodity rather than to go without (i.e. they are encouraged to 'trade at the margin') and likewise for the sellers.

Each experiment is run as a sequence of distinct trading sessions (or periods or 'days'). At the start of each trading day, all the traders are allowed to make verbal 'shouts' i.e. quotes of a price; sellers shout offers (e.g. "sell at $2.5") and buyers shout bids (e.g. "buy at $1.20"). The shouts continue with both groups of traders altering their shout-values in order to secure a deal. In a typical experiment, trading in the first day is characterised by early transactions taking place at prices differing significantly from the equilibrium value (obtained as the intersection of the demand and supply curves which arise from the distribution of limit prices for the traders). As the day progresses, the transaction prices approach the equilibrium. On subsequent days, this convergence is achieved faster. Smith introduced a "coefficient of convergence", α (Alpha) that is computed at the end of each trading day. The metric is defined as below

$$\alpha = 100\ \alpha_0/P_0$$

13 Incidentally Vernon Smith won the Nobel Prize for Economics in 2002 for his work on experimental economics.

where α_0 is the standard deviation of transaction prices around the equilibrium price $\mathbf{P_0}$. Smith also measured the efficiency of the market, which was observed to be close to 100% with human traders.

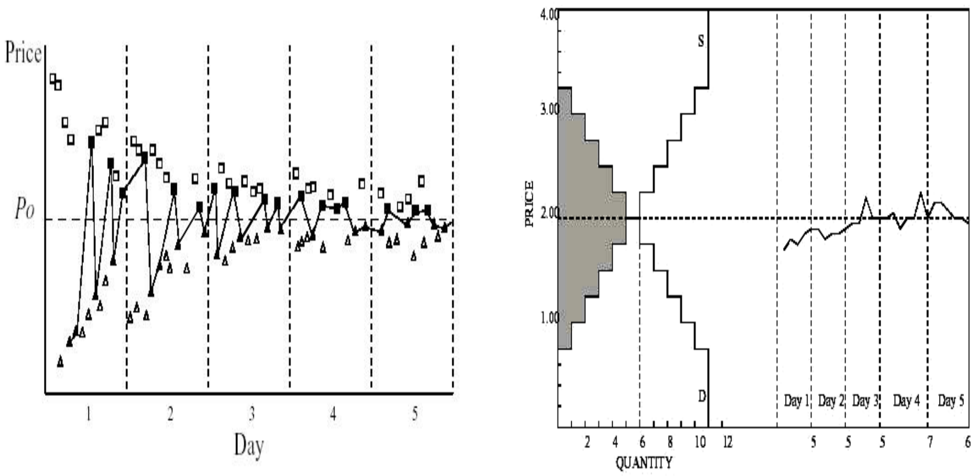

Fig 3.2 Plots from Smith's experimental CDA market with human traders. The figure on the left shows transaction prices during a 'day' while the one on the right gives the demand/supply curve used and a plot of prices over the experiment period.

Smith was able to prove the robustness of the double suction mechanism under a wide variety of supply and demand configurations, restrictions on the number of agents and conditions regulating communication between sellers. He summarises his results as follows:

"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 didn't have to have large numbers. Economic agents do not have to have perfect knowledge of supply and demand. You do not need price-taking behavior – everyone in the double oral auction is as much a price maker as a price taker." [Smith 1982]

Human beings are notoriously smart creatures: the question of just how much intelligence is required by a software agent to achieve the performance obtained with humans in the double auction markets as identified above is an intriguing one. The early literature in this field was concerned with exactly this question in the context of electronic double auction markets. The next section is a survey of the research into automated (or electronic) double auction markets, which employ software programs or agents in place of human traders.

3.2 Trading Strategies in Agent Based Continuous Double Auction Markets

This section analyses some of the literature in the field of automated markets, essentially focusing on the double auction as the mechanism and the issue of whether ‘intelligence’ in agents is necessary for trading at equilibrium prices or whether the discipline imposed by the market mechanism ensures the same.

Gode and Sunder [Gode and Sunder 1993] conducted a set of experiments similar in style to Smith’s, but which utilised “zero intelligence” (ZI) programs that submit random bids and offers to replace human traders in electronic double-auction markets. They explored the performance of both ‘unconstrained' (i.e. traders that can enter into loss-making deals by bidding below their limit prices) and ‘constrained' (i.e. traders that do not enter loss making bids or offers) zero-intelligence traders, ZI-U and ZI-C and compared the results of these traders to results from human traders operating in (almost) identical experimental conditions.

The experiments with both types of ZI traders were conducted using minor simplifications of the NYSE continuous double auction, with a transaction canceling any unaccepted bids and offers. The traders dealt in lot-sizes of a single unit of commodity. To accommodate the lack of intelligence of the traders, a deal was made whenever a bid and offer crossed: whenever a buyer made a bid higher than the current lowest offer, or whenever a seller made an offer lower than the current highest bid. In both cases, the transaction price is the earlier of the two shouts.

Differences in performance between the ZI-U and ZI-C traders, and between the ZI-C and human traders, could indicate the different extents to which overall market behavior is dependent on human intelligence or market structure:

“*The difference between the performance of the human markets and that of the ZI-C markets is attributable to systematic characteristics of human traders. If ZI-C traders are considered to have zero rationality, this difference in performance would be a measure of the contribution of human rationality to market performance. On the other hand, the difference between the performance of markets that do impose a budget constraint on ZI traders and the performance of those that do not is* *attributable to the market discipline. 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 outcomes to the discipline imposed by the double auction on traders*" [Gode and Sunder 1993]

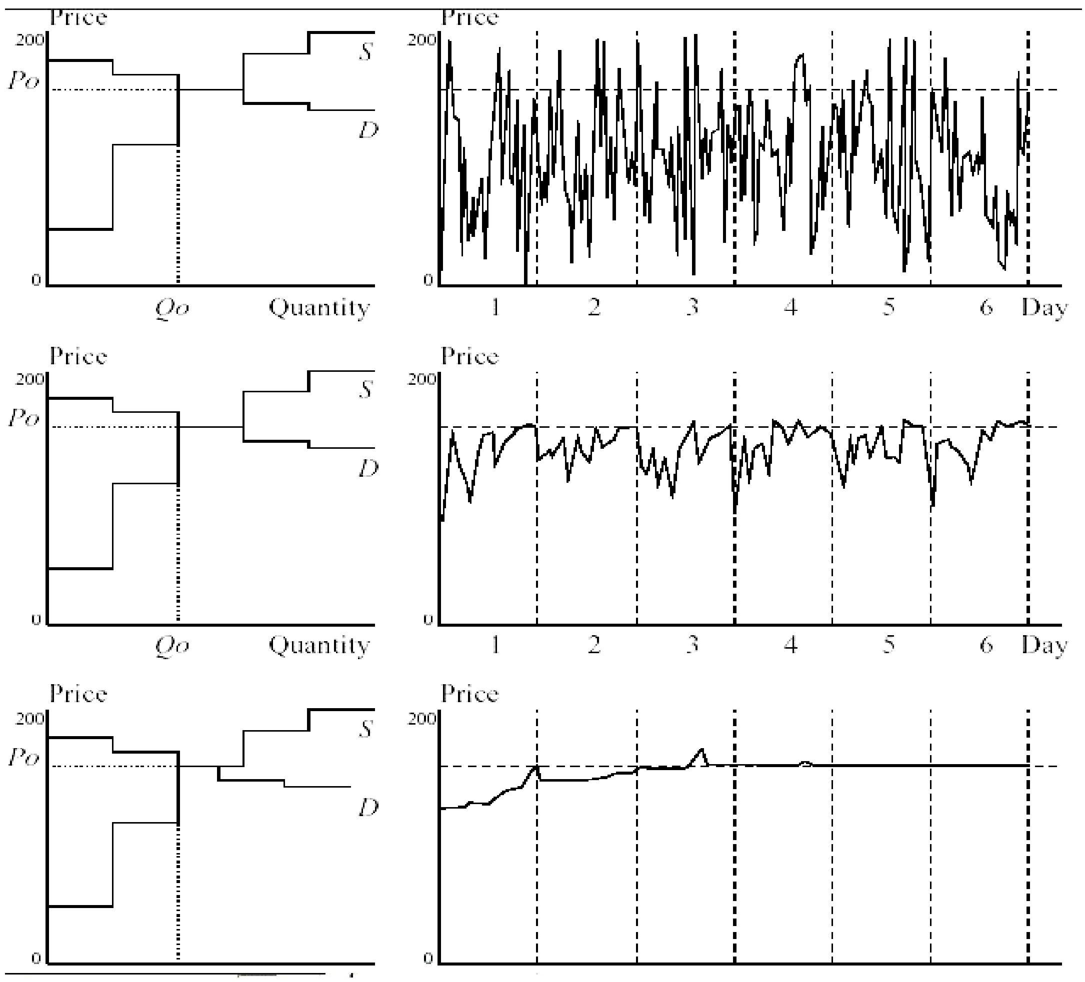

Fig. 3.3 Plots from Gode and Sunder's experiments. Top: results with ZI-U traders. Middle: Results with ZI-C traders. Bottom: results with human traders. The figures on the left are the demand and supply curves used.

In a critique of the above study, Cliff [Cliff 1997; Cliff and Bruten 1997] indicated that Gode and Sunder's results were a consequence of their experimental regime i.e. a result of the gradient of the demand & supply curves being roughly equal (for details see Cliff [Cliff 1997; Cliff and Bruten 1997]). He introduced the notion of Zero Intelligence Plus traders or ZIP traders that are simple agents, which make stochastic bids (like ZI agents). However these agents also employ an elementary form of learning, which allows the agents to adjust a given profit margin dynamically. Cliff was able to prove that the ZIP agents achieved robust convergence in CDA markets in general with differing demand supply conditions.

- For SELLERS:
  - if (the last shout was accepted at price q)
  - then
    1. any seller s_i for which $p_i \leq q$ should raise its profit margin
    2. if (the last shout was a bid)
       then
       1. any active seller s_i for which $p_i \geq q$ should lower its margin
  - else
    1. if (the last shout was an offer)
       then
       1. any active seller s_i for which $p_i \geq q$ should lower its margin
- For BUYERS:
  - if (the last shout was accepted at price q)
  - then
    1. any buyer b_i for which $p_i \geq q$ should raise its profit margin
    2. if (the last shout was an offer)
       then
       1. any active buyer b_i for which $p_i \leq q$ should lower its margin
  - else
    1. if (the last shout was a bid)
       then
       1. any active buyer b_i for which $p_i \leq q$ should lower its margin

Fig. 3.4 A pseudo-code representation of Cliff's ZIP algorithm [Cliff 1997]

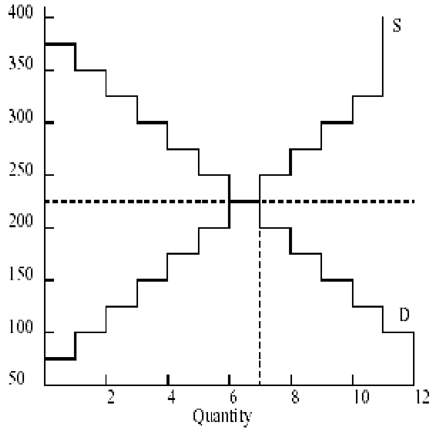

Figure 50: Supply and demand for market where only sellers shout: 12 buyers and 11 sellers. Theoretical equilibrium price $P_0 = 2.25$; quantity Q_0=7.

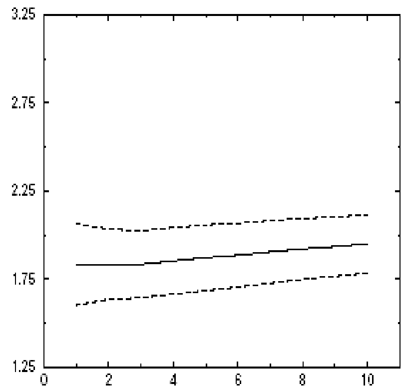

Figure 51: Mean ZIP transaction prices, averaged over 50 experiments, for the market of Figure 50. Format as for Figure 25

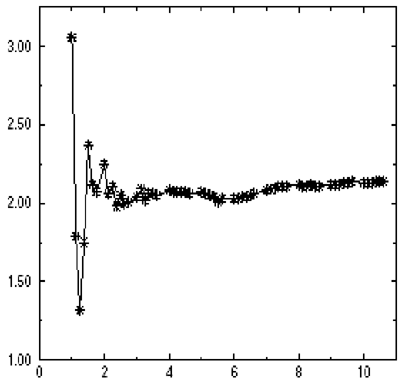

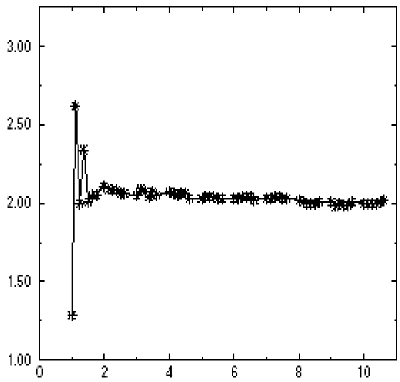

Fig. 3.5 Experimental results from Cliff's[Cliff 1997]experiments with ZIP traders. The figures on the top are explained in the accompanying text. The ones below are plots of transaction prices for individual experiments with equilibrium price as 2 units.

These results serve to establish a kind of minimum intelligence level required for trading at equilibrium levels in continuous market environments like the CDA. However Cliff's experimental setup did not quite model a CDA as it occurs in real world markets. It was designed for simulated markets with no explicit notion of time, no persistent orders (i.e. orders were allowed to lapse if not traded within a specified duration), and no definite end of a trading period. The mechanism involved randomly selecting a buyer or seller to submit an order and then polling the other side of the market to see if anyone was willing to accept it. If so, a trade takes place, otherwise a failure was recorded. A trading period came to an end after a specified number of failures are recorded. Priest et al [Priest and Tol 1998] and later Das et al [Das, Hanson et al. 2001; Tesauro and Das 2001] independently provided a modified form of Cliff's algorithm in a more realistic Continuous Double Auction scenario.

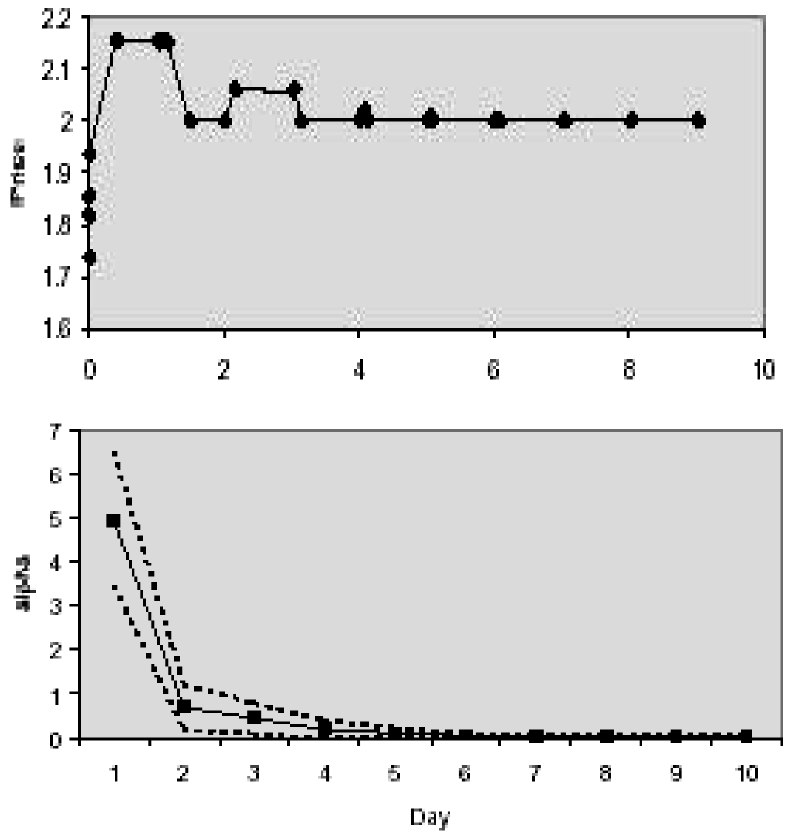

Fig. 3.6 Results from Priest et al [Priest and Tol 1998] The figure on the top shows transaction prices in an experiment with theoretical equilibrium price as 2 units. The bottom figure shows the change in Alpha during an experiment.

Gjerstad and Dickhaut [Gjerstad and Dickhaut 1998] introduced a more sophisticated algorithm for agents trading in the CDA. The use a belief based modeling approach to generate appropriate bids in an auction. An agent constructs an order and trade history *H*, consisting of all orders and trades occurring since the earliest order contributing to the M_{th} most recent trade. From the history *H*, a buyer or seller agent forms a subjective "belief" function *f(p)* that represents its estimated probability for a bid or ask at a price p to be accepted; e.g. for a seller,

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

Where *AAG(p)* is the number of accepted asks in *H* with price ≥ *p*, *BG(p)* is the number of bids in *H* ≥ *p*, and *UAL(p)* is the number of unaccepted asks in *H* with price ≤ *p*. Interpolation is used to provide values for *f(p)* for prices at which no orders or trades are registered in *H.*

The GD agent [Gjerstad and Dickhaut 1998] then chooses a price that maximizes its expected surplus, defined as the product of *f(p)* and the gain from trade at that price (equal to $p - l$ for sellers and $l - p$ for buyers, where l is the seller cost or buyer value). Thus the algorithm does not require the knowledge or estimation of other agents' costs or valuations. A comparative assessment of bidding strategies in double auctions is undertaken in [Rust, Miller et al. 1992; Tesauro and Das 2001]

Other works include Vulkan et al [Vulkan and Priest 1999] which proposes a learning mechanism that combines belief based learning with reinforcement learning. Park, Durfee et al [Park, Durfee et al. 1999] propose a bidding strategy (called the p-strategy) based on the idea of agents building a stochastic model of the Double auction market process using Markov Chains. He, Leung et al [He, Leung et al. 2002] have developed a bidding strategy for CDA's based on heuristic fuzzy rules and fuzzy reasoning mechanisms and compared it with the other prominent strategies described here.

The work of Das et al [Das, Hanson et al. 2001; Tesauro and Das 2001] is interesting in that it matches agents with artificial bidding strategies against human traders in computerized double auction markets, and shows that simple bidding strategies described above can overcome inexperienced human traders.

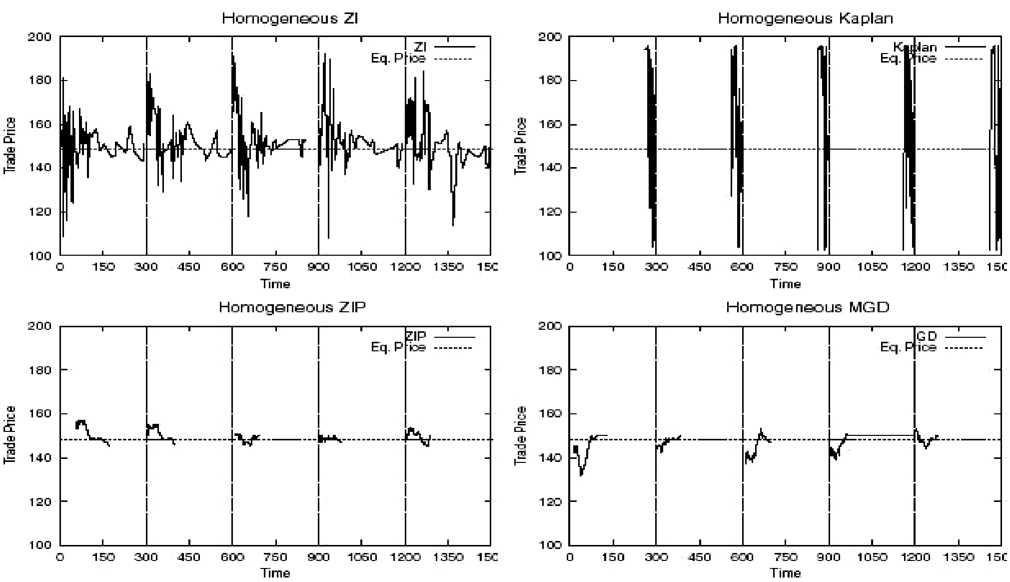

Fig. 3.7 Results from Tesuaro and Das[Tesauro and Das 2001] who compare different adaptation strategies for a CDA market.

3.3 Quote Driven Market

The Internet has served as a fertile ground for the emergence of a wide variety of electronic marketplaces, which can be thought of as the precursors of automated markets of the future. Online exchanges for chemicals, electronic components, forest products, metals, energy, bandwidth and even flowers exist. The characteristic features of successful electronic marketplaces have been cited as liquidity, integrity, efficiency and flexibility [Kollock 2000]. This section discusses the role of liquidity in the context of electronic markets and endeavours to develop a case for an alternative model to the auction based markets currently adopted as the mechanism in most electronic marketplaces. The Quote Driven market has been widely implemented and studied as an alternative to the auction based (or order driven) institutions in the financial markets world [Madhavan 1992]. The field of economics, which studies the structure of financial markets and the transformation of client orders into trades, is known as *market microstructure*.

3.3.1 Market Microstructure and the organization of Financial Markets

The essential difference between the analysis of financial assets markets and that of the simple commodity markets studied in this chapter relates to the type of goods being traded. Financial assets like securities, futures, bonds etc, are typically long lived rather than having value for a single period. They derive their value not just from the current sales or valuation, but also from a stream of dividends that accrue over time. In other words, the simple commodity markets discussed earlier involved private value considerations, whereas financial markets deal in instruments, which have public or correlated values. In such a case, uncertainty becomes a problem because the current value of an asset depends upon the expectations regarding resale prices or future dividend values.

Market Microstructure [Cohen, Maier et al. 1986; O'Hara 1995; Madhavan 2000] is the study of the processes and outcomes of exchanging assets under explicit trading rules. Much of economic literature abstracts from the mechanics of trading, whereas microstructure literature specifically analyses the effect of trading mechanisms on the price formation process. In economic theory, the price of a commodity is determined by the supply and demand curves for the particular item. In practice however, the prices need not equal the full information expectations of value because of a variety of frictions, which introduce additional costs. These costs arise due to the uncertainty and asymmetric nature of the information available to the market participants, as decision and transaction costs. Essentially, prior to the development of market microstructure the major body of literature in economics assumed the behaviour of competitive market prices to be like those in a *Walrasian* auction (an analogy could be the study of Newtonian mechanics in physics which considers point masses in vacuum). In such a scenario there is the implicit assumption that there are no costs associated with the process of trading. Such costs may arise due to charges levied by a particular market, or more importantly they may reflect the implicit costs associated with immediate trade execution. This is a reasonable assumption because the trading process generally has a time dimension (unlike the *Walrasian tatonnement*). Thus, over time, the number of buyers and sellers might be equal, but at any particular time this is not guaranteed. In such a case, the imbalance of trade would make it impossible to find a market clearing price at a given time *t*. The earliest investigations into market microstructure studied the operation of a special class of market participants (dealers/market makers) who provide immediacy by standing ready to buy and sell at stated prices (i.e. the bid-ask prices). The bid-ask spread (difference between the bid and ask prices) quoted by the market maker was deemed to be the ‘cost of immediacy’ or the return per share to the providers of immediacy (i.e. the market makers).

Market microstructure has of late acquired a much broader field of interest. Here the discussion that follows will be limited to a focus on informational issues dealt with in this field, which are relevant for our purpose. The importance of information in decision making; i.e. agent’s behaviour and hence market outcomes are very sensitive to the assumed information structure. Madhavan [Madhavan 2000] identifies some of the main categories of research in this field as under:

1) *Price formation and price discovery*: A look inside the “black box” which helps explain how customer demands are translated into realised prices and volumes.
2) *Market structure and design issues*: Focuses on how different rules affect the black box and hence liquidity and market quality.

3) *Information and disclosure*, especially market transparency: Discusses how revelation about the workings of the black box affects the behaviour of agents and their strategies.

3.3.2 The Quote Driven Market model

Figure 3.8 below gives a generic representation of a Quote Driven market [Bourne 2000; Bourne and Zaidi 2001].

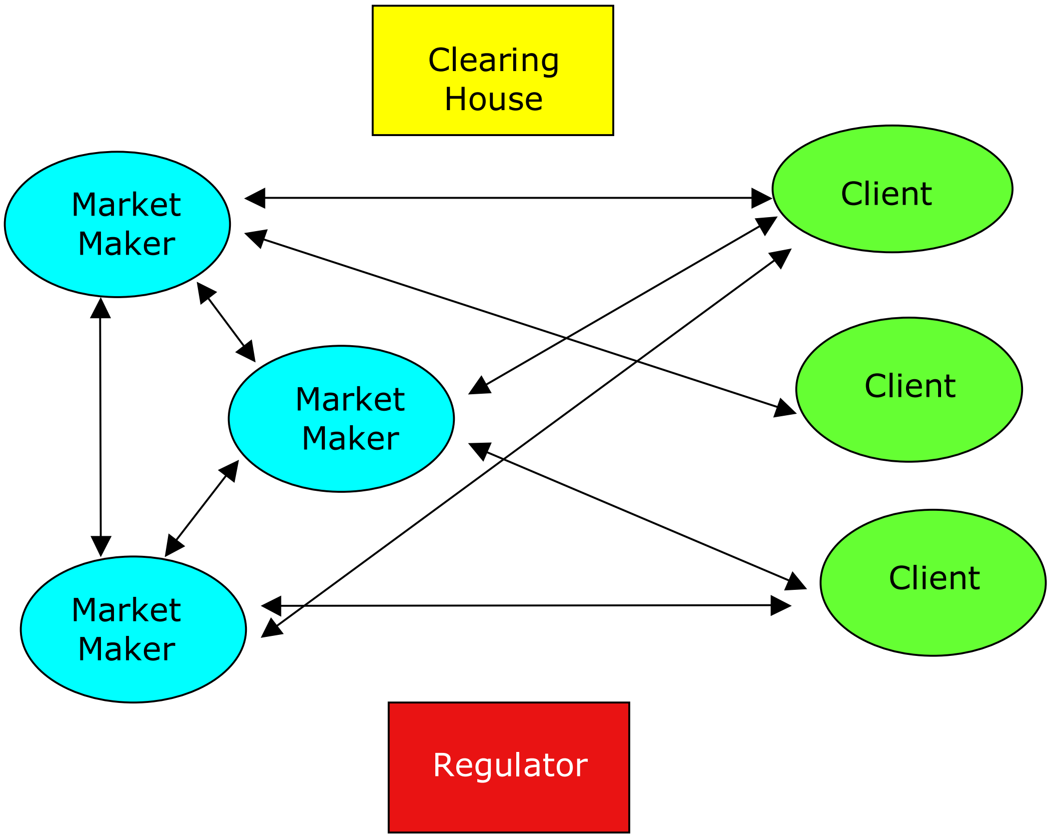

Market makers
- quote two way prices
- bid (price willing to buy for)
- offer (price willing to sell for)
- act as counterparts in interactions
- have certain obligations to trade

Clients
-buy and sell through Market makers

Fig. 3.8 Structure of Quote Based Market

The essence of the Quote Driven market lies in the presence of market makers (or traders) that act as intermediaries between the clients (i.e. the buyers/sellers). The market makers act as principals in transactions, taking on positions in the product being traded; a trader with a long position has product to sell whereas one with a short position has sold more than he owned. The traders thus provide a service to clients: each client is free to buy or sell product whether or not there exists a willing counterpart in the marketplace; the trader provides the price and acts as the counterpart in the transaction. Clearly, traders expect to benefit financially from providing this service and, conversely, clients must be willing beneficial to pay some amount for the service provided. One of the aims of this research is to determine what a fair price for such a service should be, and to what extent this will depend on the market parameters.

In a Quote Driven market, a client wishing to buy or sell quantities of product requests live prices (or quotes) from traders, one of whom subsequently acts as the counterpart in the client's transaction, should his price be acceptable.

Note that the traders do not (necessarily) have any end use for the product being traded; they purely act as go-betweens in the market. This contrasts both with order-driven markets, in which counterparts are matched against each other, the traders taking a commission on the transaction, and with auction protocols, in which buyers and sellers come together synchronously to trade, requiring a third party (the auctioneer) to determine the winner of the auction and the price at which goods are to be exchanged. Thus, in a Quote Driven market, fluctuations in the client's supplies and demands are smoothed over by the traders; the market as a whole stands to benefit from this buffering effect, which ought to lead to greater liquidity and less price volatility.

The Quote Driven market operates through a system of clearly defined market parameters and roles, decentralized real-time trading and centralised but longer-term enforcement of regulations by the institutional body, to which all clients and traders belong. Clients are free to trade with any trader and interact with them by requesting quotations. Each trader must make a live, two-way price at which he is obliged to trade with the client in any quantity within the pre-set market trading limits (the market size). The quote is composed of a bid price, at which the trader is willing to buy, and an offer price, at which he is willing to sell. The maximum difference between the bid and offer is fixed and known as the trading spread. Thus whenever a trader quotes a price, he does not know which way the client is inclined to deal and therefore has an in-built incentive to quote a price that reflects his true valuation of the current market value of the product. Additionally, traders may request quotes from, and trade with, other traders in amounts up to the market size. This allows traders to square their positions by selling a long position or covering a short position, and to take advantage of any differential pricing. Moreover, the competition among the traders for the incoming client trade forces them to reduce the bid-ask spreads so that eventually in a perfectly competitive environment their profits reduce to zero (this is however rarely the case in practice; spreads do however tend to reduce in the face of aggressive competition).

The automation of communication (i.e. request-quote-trade-acknowledge) in the above setting does not require any sophisticated technology other than secure authentication protocols. However, modeling the behaviour of the traders, the prices made and interaction with other traders is a complex process. The next section considers the market parameters in a Quote Driven market following which we discuss some relevant issues regarding the behaviour of the clients and market makers with the help of theoretical models in such settings.

3.3.2a Market Parameters

The operation of the Quote Driven market is determined by several key market parameters:

1) *Trading (or bid-ask) spread*: The difference between the market makers bid and offer price. In a stable market, the trader will make a positive profit by buying on his bid and selling on his offer i.e. the spread acts like a commission. However the spread can also be thought of as providing a buffer to the market makers against adverse movement in the market, which given the trader's position in the market will affect his profits.
2) *Market trading limits*: The maximum and minimum quantities of a product for which a trader's price is valid. Can be specified as discrete limits over which the market makers prices remain valid. Usually larger maximum sizes correspond to a more liquid market.
3) *Number of traders*: This represents a critical issue in the efficiency of a Quote Driven market. Essentially, large numbers of trader's would imply greater competition (thereby reducing the bid-ask spread) whereas a single trader will effectively act as a monopolist, providing a poor service for the clients.
4) *Response times*: Two time limits can be identified. Firstly, the traders should reply to the request for quotes within a fixed time interval and secondly, the

quotes are valid for limited periods only and clients should deal before the expiry of that period.

5) *Cost of Running Positions*: There are definite costs to the trader for maintaining a position on either side of the market. There might be borrowing costs associated with being short as also with running long positions (e.g. owning bandwidth that remains unused over long periods may be costly).

3.3.2b Price Formation in Quote Driven markets

Much of the early literature in market microstructure concerns the operation of market makers in financial markets. By virtue of their central position and role as price setters, market makers provide the logical starting point in the exploration of how prices are actually determined in the market. As mentioned previously, the earliest models of market makers viewed them as providers of liquidity who were compensated for their services by the bid-ask spread. The market maker has a passive role i.e. simply adjusting the bid ask spread in response to the changing conditions.

3.3.2c Behaviour of the Market Maker:

The Role of Inventory

The early view of market makers as passive providers of liquidity was modified by studies, which suggested that the market makers actively adjusted the bid-ask spread in response to their inventory positions. The primary role of the market maker can still be considered as a provider of immediacy, but he also takes an active role in price setting primarily with the objective of achieving rapid inventory turnover and not accumulating significant positions on either side of the market. The idea is, that as the dealer trades, the actual and desired inventory positions diverge forcing the dealer to adjust the prices. Since this may result in unexpected losses, inventory control implies the existence of a bid-ask spread even if the actual transaction costs (physical costs of trading) are negligible.

Models of market maker inventory control [Garman 1976; Stoll 1978] typically use stochastic dynamic programming. Essentially these models envision the market maker facing a series of mini auctions through the day. As the number of trading rounds become arbitrarily large, the trading process approximates that of a continuous double auction. Inventory models provide an added rationale for reliance on market makers. Just as physical markets bring together buyers and sellers in *space*, the market maker does so in *time* through the use of inventory. The presence of market makers who carry inventories imparts stability to the price movements through their actions relative to an automated system that simply clears the market at each auction without accumulating inventory.

Although market microstructure provides a further model of market maker behaviour[14], we will restrict our focus to inventory led pricing mechanisms here. This is due to the fact that we are concerned with simple commodity markets and a comparison with the Double auction based markets in such settings.

3.4 Applications

Automated markets for the trading of goods in open distributed e-marketplaces are still a long way off. Current research efforts will hopefully lead to the development of robust and sophisticated automated price negotiation in such domains. A glimmer of the future possibilities can be had from simple implementations such as the *elves* or automated bidding agents offered at auction sites such as eBay and Amazon and by third party services such as eSnipe (eSnipe automates a common practice among eBay bidders of waiting until a few seconds before an auction's close to place a bid), pricebots such as buy.com which automatically undercut the competition and shopbots such as DealTime that minimise the total cost of a bundle of goods by partitioning it across one or more vendors, taking shipping cost schedules into account.

Although many markets including travel, music and book markets have undergone dramatic changes due to the development of e-commerce, it is the utility markets in gas, electricity and bandwidth etc. that hold huge potential for major development. Much of the change has been brought about as a result of restructuring within these industries which has seen the introduction of competition (in the provisioning of these services to the end user) where earlier there used to be monopoly control by single organisations. We focus here on the emerging electricity

[14] The Asymmetric Information based models ([45 49]) are utilised in the study of financial markets where the commodity that is traded are derivatives (stocks/bonds/futures etc.) Such goods are much more complex than simple commodities in their representation and have extra considerations.

and bandwidth markets and indicate the possibilities that exist for automating interaction with agent technology.

3.4.1 The Deregulated Electricity Markets

In many countries, power markets have been opened up for competition during the last decade or so (e.g. UK in 1989, Norway, US, Australia etc. See [Wolak 1997; Chao and Wilson 1999; Ygge 1999]). The initial term used for the kind of restructuring that followed was deregulation, however later it seemed sensible to adopt the term re-regulation instead, as the new markets were at least as regulated as the previous ones. The legal aspects of these changes effectively meant that (among other things):

- The electricity sales and the electricity production were separated from the transmission and distribution of electricity.
- The distribution stayed a monopoly, while the sales and production went competitive.
- The distribution companies are obliged to transmit power, sold by any trading company, to their customers.
- All customers have the right to change electricity supplier.

Besides the legal aspects, structural changes meant the introduction of radical designs as seen in the inception of electricity trading pools (the Nordic Pool), exchanges (as in California, England & Wales after NETA) and bilateral markets. In general, Figure 3.9 below provides a simple representation of this change.

In the traditional regulated power market (top half of figure 3.9), the customer has no possibility of choosing a provider. There was a single power utility responsible for both production and transportation. In the re-regulated market (bottom half of figure 3.9), most of the power is sold through a power pool (or exchange), though bilateral contacts are also allowed. At the moment the power pool serves resellers of power and very large customers. Transportation of the electricity is separated from production both at the provider and customer sides. One important aspect of trade in electricity markets is the obvious concerns to do with grid or system stability. As such, the network operation is managed by an independent system or grid operator (ISO), which effectively manages the trade in the power pool, deciding on the various combinations of production and consumption (expressed as bids by the generators and resellers) that stabilize the grid.

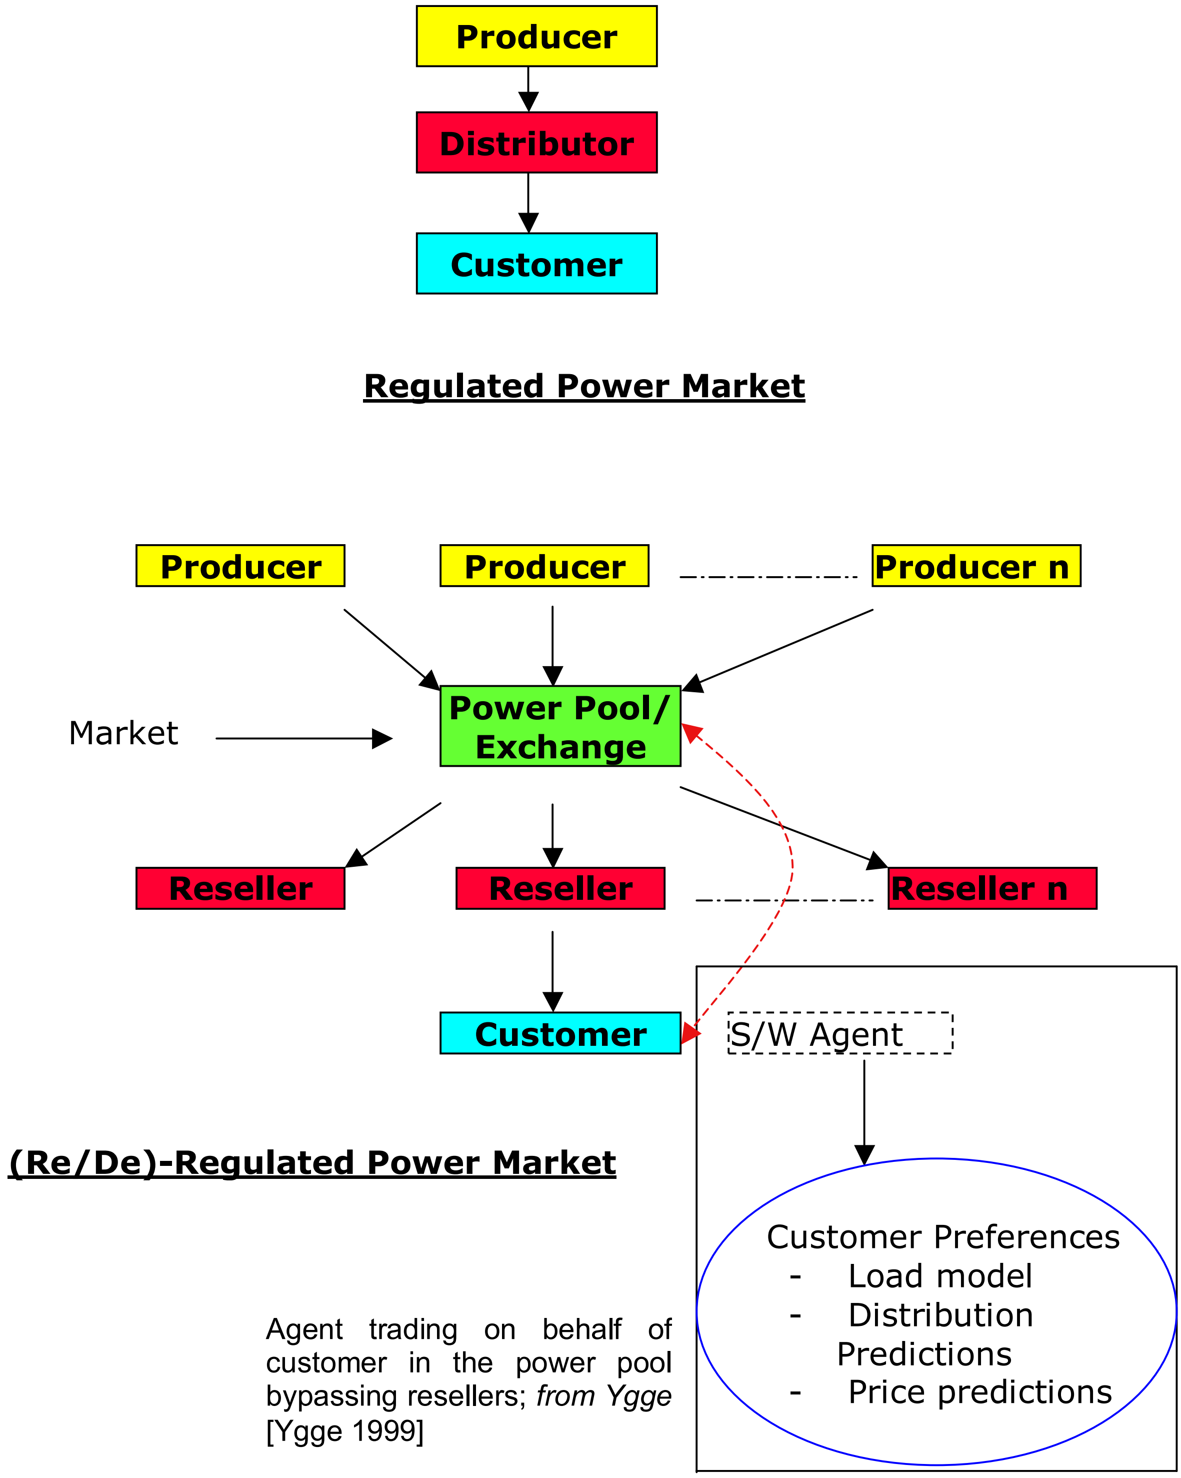

Fig. 3.9 The structural change in the power industry

The power market is generally divided into forward and spot markets. Usually, the spot market is referred to as the day-ahead market for supply of electricity for half-hourly (or hourly) intervals the following day. There is also a near real time market in most electricity markets, which is also managed by the system or grid operator. The authority exerted by the ISO in its control over the temporally differing spot markets (day ahead or almost real time) differs appreciably. The ISO has almost complete authority in real-time operations. More latitude is allowed for decentralized markets to establish plans as the time horizon increases. Besides the spot market, there is also a futures or forwards market that sells contracts, which are obligations to buy or sell at a future time a given quantity of electricity at a price agreed upon at the time the contract is entered into. These contracts are standardized along various dimensions e.g. the amount or quantity of the electric power, the time horizons to the contracts which could be week-long contracts, ranging from up to 4-7 weeks in the future; block contracts of 4 weeks long, for electricity delivered up to a year in the future; and seasonal contracts for blocks lasting an entire season of the year, for electricity delivered over one year in advance. The time period within the day may also be specified in the contract.

A significant aspect of trading in electricity markets is the degree to which a market is residual after previous transactions. Typically, a sequence consists of bilateral or organized futures/forwards markets, day-ahead or hour-ahead markets and the real time market. Often each market in the sequence is mainly a residual market for the trading deviations from transactions in previous markets. For instance, around 24 hours before the physical delivery (i.e. in the day ahead market) suppliers tend to fine tune their positions i.e. buying and selling electricity to cover excess or shortfall, between their actual position and that covered by the contracts in the forwards/futures markets.

As indicated above, presently the power pools and exchanges cater to trading between the generating companies, resellers and large customers. Ygge [Ygge 1999] presents an example of a software agent called HOMEBOT which is an extension of their software agent (originally proposed for direct power load management at customer sites) for trading in power pools (Figure 3.9 above depicts this graphically). The HOMEBOT utilises such information as customer preferences e.g. how much a user is willing to pay for maintaining certain temperature, a load model for the physical surroundings, disturbance predictions (e.g. weather) and price predictions to compute a utility function. The utility function represents the valuation of the resource (electric power) to the customer. With knowledge of the utility function, the HOMEBOT can generate a demand function for the consumption requirements of the user for (say) each hour in the next 24 hours. This demand function is declared by the HOMEBOT as a bid to the pool. The ISO then computes the market-clearing prices for each of the 24 hours and allocates the resource (electric power) amongst the participating agents (the actual allocation or physical supply occurs in real time after the ISO has balanced any deviations from the day ahead market to the real time market). Ygge points out that one advantage of utilising software agents to trade on our behalf in an automated market would be to reduce the granularity of the market from hour/half-hour durations to minutes and even seconds. The improved time resolution could help the agents submit more accurate demand functions and even allow direct trade in near real time markets (within limits of the communication capabilities of the communication infrastructure). It would also allow for increased economic effectiveness in energy trading and better handling of control and crisis situations on the electricity network. Ygge suggests the presence of a continuous market for enabling the agents to trade rapidly in real time (and online). This prospect raises the possibility of the entire electricity grid becoming unstable (because the stability of the grid is dependant on the allocations not changing too rapidly, particularly between regions). Although it is difficult to establish the effect on grid stability in the absence of further tests, it is the author's belief that the trading mechanism employed in such a continuous market will determine to a large extent the effect of such trades in real time.

3.4.2 Bandwidth Trading

The telecommunications industry is undergoing rapid change both in terms of technological advancements and the application of more efficient and economic management of available resources [Courcoubetis and Weber 2003; Economides 2004; Varian 2004]. In recent years, the unprecedented abundance of bandwidth supply[15] (in fixed networks), and the advances in switching/routing and network management, has removed some of the barriers in provisioning bandwidth on demand. This has led to the trend of suppliers and consumers actively trading excess

---

15 Rapid interconnectivity and capacity expansion has led to localised or marked wide surplus.

capacity leading to the formation of bandwidth markets. Lehr & McKnight [Lehr and McKnight 1998] identify the following reasons for the presence of bandwidth markets:

1) The competitive and innovative nature of the communications infrastructure value chain (seen in the emergence of multiple service providers; i.e. non vertically integrated firms) fuels the need for robust wholesale markets for the underlying transport services.
2) Rapid network interconnectivity and capacity expansion that results in localised and/or market wide surplus. The owners of this surplus have the incentive to lease the excess capacity.

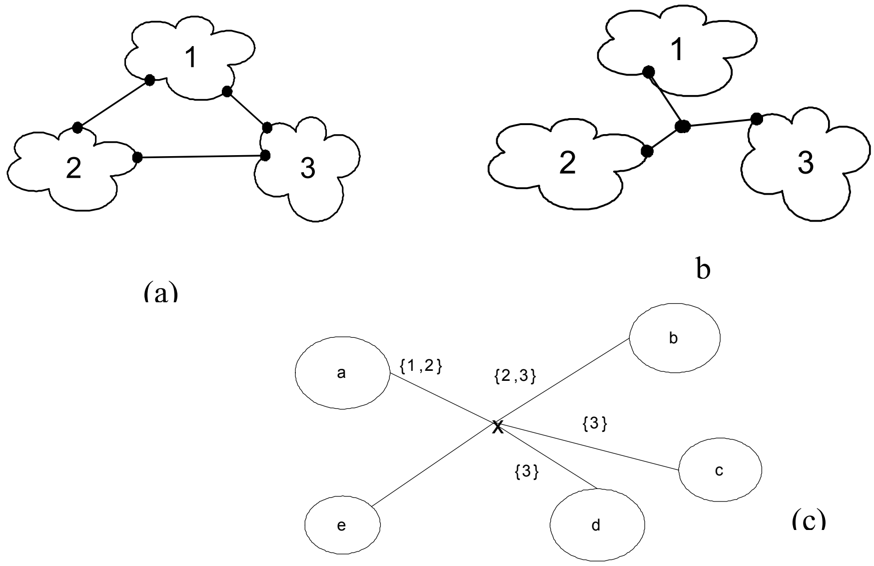

Fig. 3.10 Expanding the network's reach (a) private peering points (b) multilateral peering and (c) commodity markets

3) The demand uncertainty for Internet traffic, which creates difficulty in network provisioning and facilities planning.
4) The underlying packet-switched network architecture offers more routing options or substitutes than a circuit switched network. This helps drive the need for bandwidth markets by potentially supplying more depth or liquidity.

In common parlance bandwidth trading refers to the wholesale trading of bandwidth or related commodities as above. (See www.Band-X.com, www.Arbinet.com etc.) This includes circuits between city pairs, Internet access, circuit switched telephony minutes, spectrum capacity[16] [Grigonis 2001; Mindel and Sirbu 2001] etc. End users are not affected by the cost and as such prices affect the traffic on a very coarse scale. A more efficient market for bandwidth allocation could be visualized as the fine-grained negotiation about access to scarce resources that takes places between end users. Most work on bandwidth markets at this lower scale is concerned with improving network performance by providing admission control at the edges [Gibbens and Kelly 1999]. A user is not admitted into the network if the network does not provide sufficient service quality. Kelly [Kelly 2001] models interconnections as reservations along a specified path and establishes that the system state stabilizes asymptotically, as end users change their demand according to network load. Courcoubetis [Courcoubetis, Dimakis et al. 2001] address the question of conditions in which a best effort network supports specified QOS connections for specific durations. The developed models are extended for dynamic pricing over the lifetime of accepted calls. Semret et al [Semret and Lazar 1999] model admission control to a network over an exponential distributed number of minutes as a number of n^{th} price auctions on 1 minute time slot access and prices an access option as the sum of call options on each of the time slots. Lukose et al [Lukose and Huberman 1997] use a CAPM like model to construct a mixed portfolio of network access with different service classes, in order to reduce the latency variance and mean. Cheliotis [Cheliotis 1999; Cheliotis 2000] describes a market based end-to-end IP service model with QOS based on commoditised bandwidth contracts at multiple public marketplaces. He introduces the notion of brokers i.e. middlemen who process consumer requests and compute end to end paths meeting the customer's requirements at the lowest possible costs. The broker does this in consultation with Market Data Providers (MDPs) that provide current market data regarding traded bandwidth segments from all the bandwidth exchanges (M).

---

16 In future additional bandwidth granularities including IP QoS flows can possibly be traded.

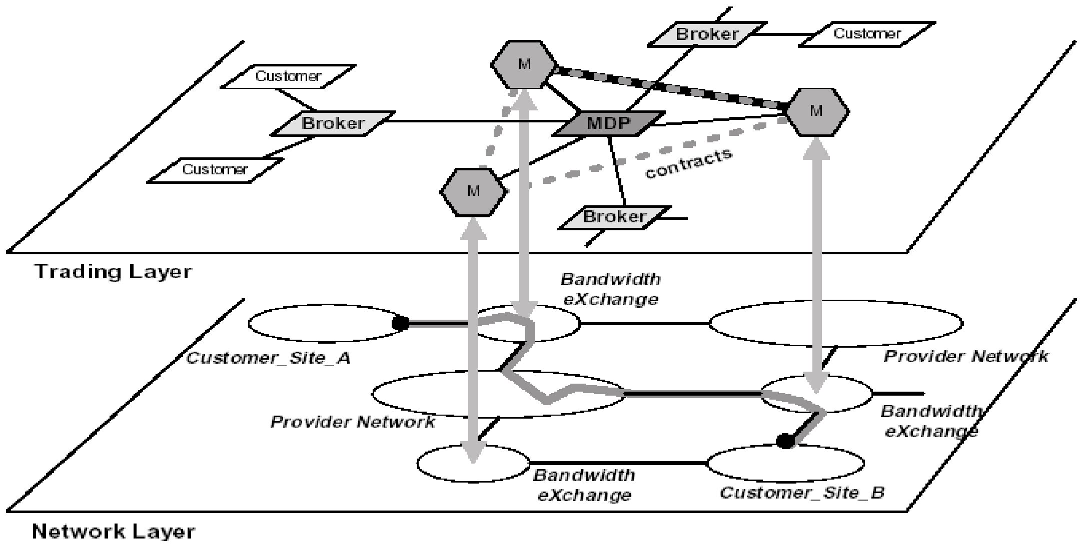

Fig. 3.11 Market based bandwidth model (from [Cheliotis 1999])

It is interesting to note that Cheliotis assumes that there exists some price setting mechanism in place at the marketplaces that matches bids and offers (he cites a double auction as a probable choice as in [Sandholm 2000]. Rasmusson [Rasmusson 2001] prices complex network services as financial derivative contracts based on the spot price of the capacity in individual routers. By choosing to trade capacity shares rather than time slot access, there is only one market per router rather than one per router per minute.

Gibney at al. [Gibney, Jennings et al. 1999] describe a market-based approach to call routing in a network, which provides the same performance measure as static call routing. The architecture describes agents that represent underlying network resources e.g. *links* and *paths*. These are traded in *link* and *path markets* (see **Fig**. 3.12) with the underlying agents using demand and supply information in the markets to adjust their prices. The link market is assumed to be a 'sealed bid double auction' (essentially a call auction) whereas the path auction is a sealed bid auction where the path agents bid to provide source to destination connectivity to the call agent (i.e. the caller).

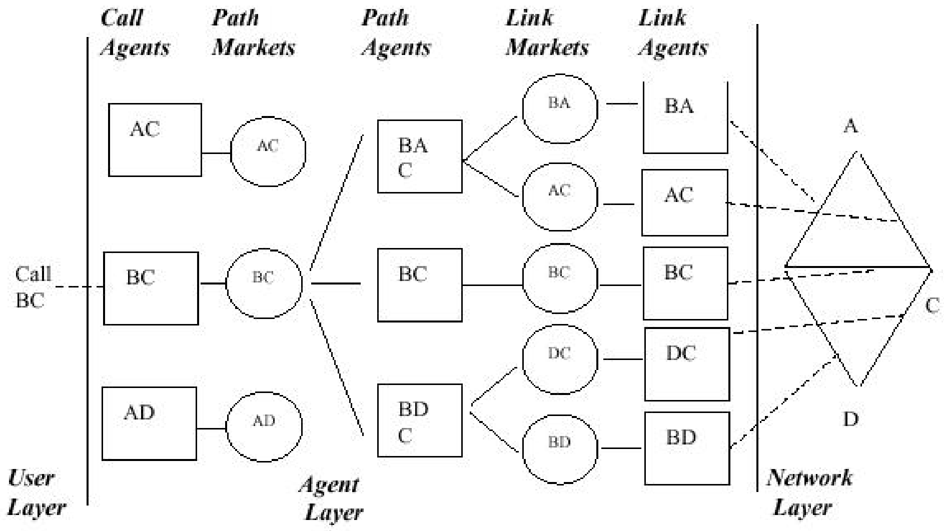

Fig. 3.12 System Architecture *from* [Gibney, Jennings et al. 1999]

In other related agent based CAC schemes, Bigham et al [Bigham, Cuthbert et al. 1999] describe a distributed architecture where the underlying resources (source to destination virtual paths of possibly multiple network/service providers) are managed by *resource agents (RAs)* that participate in auctions conducted on behalf of a user (that requests a connection with defined QOS) by a *connection agent (CA)* . The authors consider a private value first price sealed bid auction for winner (the RA that wins the auction) determination.

3.4.3 A New Service Model for the Telecommunications Industry?

The expense of deploying Third generation (3G) Telecommunication Systems will be huge. This is evident from the billions spent by operators (about £22 billion in UK alone) in acquiring the spectrum rights, with equipment outlays expected to cost in the same region. With cheap (core) network bandwidth and the highly competitive market driving bandwidth costs down, many an operator around the world has been plunged into a financial crisis. This has lead to a radical rethinking where the monolithic and highly integrated telecommunications architecture provided by the likes of BT and AT&T could be replaced by a dynamic confederation of multiple (competing and sometimes cooperating) service providers [Katz, Stoica et al. 2002].

For instance, the current approach to operating wireless communication systems is deemed outdated and inefficient on account of the below:

- Scarce spectrum resources are statically partitioned among licence holders, independent of subscriber density or nature of the users workloads. A much more efficient situation could be that an operator with under-utilised spectrum resources resells capacity to an operator which is oversubscribed on a short term, as needed, basis. This can be seen to a limited extent in the mushrooming of MVNOs (Mobile Virtual Network Operators) [Mitchell and Moore 2001] (e.g. the relationship between Virgin Mobile (a MVNO) and T-Mobile in UK) and more recently the agreement between T-Mobile and MMO2 to share 3G infrastructure (Various OFTEL[17] publications [OFTEL 2002]). However a much higher degree of efficiency is obtainable through a peer-to-peer model rather than the hierarchical model of MVNOs as above.
- Cell sites for antenna deployment currently lead to duplicate coverage across providers. With antenna sites becoming scarce due to local resistance and other factors, a much better way would be to have antenna operators, which provide a resource (antenna access) to subscribers thereby removing the need for the current duplicative coverage efforts.

The SAHARA project[18] envisages an emerging architecture characterised (among other things) by:

- Confederated Services: The architecture would support overlapping service provider regions with subscribers able to roam among them for service provisioning. The architecture would support co-opting among service providers to share resources like spectrum, bandwidth, storage, antenna sites etc.
- Efficient allocation of scarce resources using dynamic market driven mechanisms. Trusted third parties (Clearinghouses, Resource brokers, B2B i.e. Business to Business Exchanges etc.) managing resource marketplaces in a fair, audited and verifiable basis.

---

17 Since 2004, the responsibilities of OFTEL have been passed over to OFCOM, the regulator set up to replace previously existing regulators for the UK Communications industry.

18 http://sahara.cs.berkeley.edu/

Such an open service resource allocation model would require service providers being able to advertise for and husband resources for near peak time needs. A crucial question raised in the SAHARA project is to the amount of information about current resource availability being made available by service providers. A possible solution offered therein is the through intermediary market-makers who can possibly hide the details of the participating provider's resource pools.

The next section provides full description of bandwidth trading via a Quote Driven market applied to multiple Service Provider network architecture. Such a market addresses many of the concerns of real time trading and issues regarding privacy and trust raised above.

3.4.4 Application of Quote Driven Market to Bandwidth Trading

Bourne [Bourne 2000] has proposed a Quote Driven market mechanism for a telecommunication network in a many-to-many scenario i.e. multiple SPs (Service Providers) compete against each other for providing services to the user (See Figure 3.13). The essence of this market is that the SPs act as market makers or brokers of bandwidth between the Network Provider (NP in the diagram) and the end user (U). The scenario of multiple service providers competing against each other is more realistic given the deregulation in the telecommunication industry.

Three types of interactions are identified in the scenario below:

1) *Network Provider/Service Provider Interaction*: This is assumed to occur over relatively longer periods.
   a) SP specifies the source-destination (s-d) pair, units of bandwidth and length of lease.
   b) NP quotes a price.
   c) SP either accepts terms and lease commences, or declines
2) *User/Service Provider Interaction*: The interaction followed is as below
   a) U requests the price for a particular s-d pair.
   b) SP quotes a price.
   c) U either accepts terms and connects, or declines.

Unlike an auction mechanism, the user dictates the number of alternative prices he sees and from which SP to obtain a quote. This is beneficial in implementing reputation mechanisms i.e. the user may keep track of SPs offering the best prices and in future request only these for a quote.

3) *Service Provider/Service Provider Interaction*:

a) SP 1 requests a quote specifying an s-d pair, units of bandwidth and length of lease.
b) SP 2 quotes a two-way price, its bid and it's ask.
c) If SP 1 wishes to deal on this price, it states whether it is buying or selling and the deal takes place.
d) SP 2 requests a quote specifying an s-d pair, units of bandwidth and length of lease.
e) SP 1 quotes a two-way price, its bid and it's ask.
f) If SP 1 wishes to deal on this price, it states whether it is buying or selling and the deal takes place

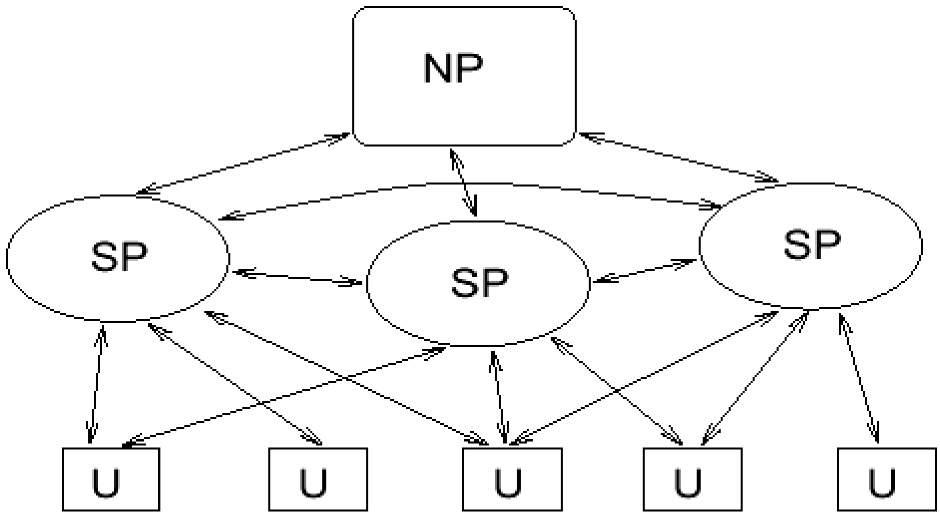

Fig. 3.13 Many-to-many market structure in the telecommunications domain

The interaction between the SPs is novel in the sense that it allows them to escape from unprofitable leases thereby providing more liquidity in the market.

Markets for telecommunication commodities can be separated into the 'spot markets' and 'reservation' or financial derivatives based markets; which is essentially the trading of excess capacity etc. The proposed model for a Quote Driven market can be utilised in both instances. However, the discussion below concentrates on a Quote Driven market model for a spot market in bandwidth trading. An automated spot market for bandwidth[19] trading will typically provide for:

- A physical interconnection point (pooling points).
- Electronic marketplace to post and accept bids and offers.
- (Independent) Quality monitoring (of SLAs)[20].
- Financial Clearing House or Credit.

At the core, the marketplace should have simple rules and provide for an *efficient* mechanism taking into account the self-interested nature of the agents. The mechanism should also be scalable to a large set of participants both in terms of computational complexity and messaging load. The Merkato® platform[21]is one such distributed and highly scalable platform which uses a proprietary mechanism; the Progressive Second Price Auction [Semret and Lazar 1999] for resource allocation in the spot market. Other auction mechanisms proposed to model the spot markets (usually CAC procedures or selection of routing paths) during connection time include sealed bid auctions [Gibney, Jennings et al. 1999] and double auctions [Cheliotis 1999].

The Quote Driven market model for bandwidth trading is different from the other models studied previously in that it captures the typical many to many relationship between Network providers (NPs), Service Providers (SPs) and the Users (U). The deregulation in the telecoms industry worldwide has led to a scenario where the users buy services from one or more service providers, who in turn buy capacity from one or more network operators. In general we can visualise the possible presence of a market at different levels (e.g. user-SP, SP-SP and SP-NP interactions). In such a scenario the user has a choice of SPs to obtain the service from. Also the NP, which owns the telecommunication infrastructure, would like to obtain maximum return by leasing out bandwidth to the SPs. The SPs in turn are exposed to the risk of leasing more bandwidth than they can sell and selling more bandwidth than they have leased.

---

19 Behind the front end systems fully automated provisioning of bandwidth does not exist at the moment but developments in switching and agent technology are leading us there.

20 There is work towards developing standardised SLAs that are machine-readable. Exchanges like Arbinet have already developed an objective view of relative QoS.

21 Invisible Hand Networks Inc. www.invisiblehand.net

3.5 Summary

This chapter discussed automated i.e. agent based trading. The discussion focussed on the experimental markets literature, specifically centring on the mechanisms that support multiple buyers and sellers e.g. the double auction. The issue of the level of 'intelligence' required of artificial entities or agents representing humans in such market scenarios is also considered. This is still a very actively researched topic in this field as the recent publicity surrounding results from IBM[22] shows

This Chapter has also indicated two promising areas for the application of agent based automated markets including electricity and bandwidth trading. We also presented an overview of the Quote Driven market and a proposed implementation to the bandwidth trading scenario which we believe is more realistic and flexible than the other models studied here. Quote Driven markets can be implemented for trading in a host of domains, including financial instruments, electricity, bandwidth and other commodities; wherever fast or real time trading may be required. Several studies have actually indicated the possibility of such market structures emerging in the new electronic commerce domain [Lehr and McKnight 1998; Rasmusson 2001].

---

22 http://www.research.ibm.com/infoecon/inthenews.html

Chapter 4

Experimental Analysis

This chapter focuses on the experimental investigation into the properties of the double auction market. It is shown that the oft-reported efficiency of the CDA does hold true in a wide variety of circumstances. However it is also borne out through experimentation that a high degree of synchronization between buyers and sellers is necessary for optimal performance in terms of efficiency and liquidity. If the buyer and seller population is varied greatly with respect to the other, the efficiency obtainable is degraded considerably.

The double auction along with the Quote Driven market forms the principal mechanisms for continuous trading environments and has been studied widely. The field of automated markets has analysed the Double Auction in considerable depth. Much of the recent research into double auctions in agent-based markets has focussed on the design of robust heuristic bidding algorithms, which are capable of functioning in real time DA environments [Cliff 1997; Gjerstad and Dickhaut 1998; Priest and Tol 1998; Das, Hanson et al. 2001] etc.

The goal of this research is not the design of high performance bidding strategies for use by the agents. Instead, the aim here is to study the market mechanisms/negotiation protocols themselves, which can be evaluated on the basis of certain criteria (e.g. Social welfare, Pareto efficiency, Individual rationality, Stability, Computational efficiency etc.) To this extent a discrete time simulator has been developed, to assess the performance of a Continuous Double auction market (CDA). Additionally, a continuous real time simulator [Bourne and Zaidi 2001] was also developed to allow for the investigation of market dynamics in a Quote Driven market comprised of a diverse population of market makers following varied but minimal learning strategies[23]. This implementation was undertaken in *Pathwalker* (a distributed agent or thread based programming environment provided by Fujitsu Labs).

---

23 The simulator was successfully used as part of an MSc coursework (IAMAS 2001-02). The students were tasked with designing their own agents to act as market makers in a Quote Driven market run by us.

4.1 Structural Properties of Markets under Minimally Intelligent Traders

Previously the motive for this research was defined as the analysis of the structural properties (i.e. the implementation of the protocols) for different market mechanisms. More specifically, the aim has been to assess through a simulation study, the market outcomes which arise as a result of the structure (or organisation) of the marketplace rather than as a result of the actions of interacting traders. This is in the spirit of the Gode & Sunder experiments [Gode and Sunder 1993] involving double auctions with Zero Intelligent traders (The ZI traders submit random bids/offers with the restriction that they are not allowed to trade at a loss). Although Gode & Sunder's ZI traders can be seen as a lower limit on the mechanistic complexity of trading agents, their results have been shown to be the consequence of experimental artefacts and do not hold generally [Cliff 1997]. Cliff [Cliff 1997; Cliff and Bruten 1997] proposed his ZIP (Zero Intelligence Plus) traders, which can be shown to achieve human like behaviour with simple learning rules. (i.e. many of the DA results can be replicated with these simplistic trader models). The ZIP traders of Cliff can be seen as an extension to Gode & Sunder's work and an attempt to identify the lower limits of intelligence observable in DA markets involving humans.

For the purposes of this study the author has implemented a minimalist behaviour on the market participants in accordance with the simple heuristics advocated by Cliff. The use of simple bidding strategies is useful in calibrating the results obtained from using different market mechanisms by restricting our focus to the effects of the market mechanism itself rather than complex strategic behaviour of the agents.

The sections below discuss the terminology and the actual experimental set-up utilised in our experiments.

4.2 CDA: Experimentation Methodology

This section summarises the experimental methodology followed by us for the Double Auction market. Tesauro & Das [Tesauro and Das 2001] give a good overview of the literature discussed here and the experimentation undertaken. While the previous work indicated here goes a long way in providing experimental verification of the properties of a DA and the limits of cognitive/learning capabilities of the software entities/agents which could be utilised in automated commodity markets of the future, it does not provide any assessment of the merits of the DA when compared to an alternative market mechanism. It is this deficiency that the author aims to address with the experimentation that follows.

The experimental methodology borrows from the earlier studies [Smith 1962; Cliff 1997; Gjerstad and Dickhaut 1998; Priest and Tol 1998; Tesauro and Das 2001] in that the demand/supply conditions in the market are induced by allocating limit prices to a predefined number of buyers and sellers in the market. The experiments are conducted in a discrete time market simulator, which allows for stochastic asynchronous bidding and trading by randomly activating a subset of agents at each time step. The solicitation of bids is achieved in random fashion. At each time step, an active agent has the opportunity to submit an order (for one unit only) or modify an existing one given certain conditions e.g. the NYSE rule[24]). When a bid and ask cross, a transaction takes place and all previous bids/asks are deleted, the traders involved drop out of trading and a fresh round of bidding starts with the remaining agents. Initially, in the first round, all participating agents make an opening bid/offer, which can be modified in subsequent rounds.

4.2.1 Definitions:

Market, *sessions*, *rounds* of bidding and *transactions* constitute a nested hierarchy in time. A *market* is divided into *sessions*, usually of fixed duration. In each session, the active buyers/sellers submit bids/asks. Each *transaction* consists of one or more *rounds* of solicitation of bids and asks from traders. At the beginning of a session, all market variables are reset and traders receive fresh endowments. Except for changes in traders' memory and wealth, all sessions are identical. For ZI traders, which lack memory and ignore wealth, each session of a market with ZI traders is statistically identical.

Each *round* of bidding culminates in a transaction or when all the active agents have made bids/offers. When bidding starts at the beginning of a round, the current bid is reset to 0 and the current ask is reset to M (an upper bound, 999 in the simulation). Bids and asks are solicited in a series of *rounds*, and submitted to a central clearinghouse. If a new bid is higher than the current bid, then it becomes the

[24] The NYSE improvement rule states that any new bid/offer must better the existing one in the market

current bid. Similarly if the new ask is less than the current ask, then it becomes the current ask. As soon as the current bid equals or exceeds the current ask, a trade occurs. The transaction price can be:

a) The current bid or current ask depending on which was submitted earlier or

b) A pre-specified weighted average of the current bid and current ask

In the simulation that follows the author has opted for choice a) for comparative purposes. If a round does not result in a transaction, more rounds follow until there is a transaction, or the session ends. The fixed duration of a CDA market session is chosen such that it gives sufficient time for the traders to transact all possible trades. Exactly how much time is sufficient to trade all possible commodities is difficult to answer in a simulated environment. For the simulations therefore, a choice has been made to limit the length of a session to three continuous rounds of bidding that do not lead to any further improvement in the bid/offer prices.

4.2.2 Sequencing

There are three possibilities with respect to the solicitation of the bids and asks in a DA [Davis and Holt 1993]. In a ***synchronized*** double auction (Clearing house) model, both the solicitation of bids/asks as well as their processing for the purpose of matching within each round is simultaneous or batched. All traders are solicited for their bids and asks. While submitting a bid or ask, a trader does not know the bids and asks submitted by other traders in that round. The highest bid and the lowest ask from the round are then used to update the current bid and current ask. If current bid exceeds or equals current ask, a transaction is completed, the current bid and current ask are reset, and the first round of bids and asks for the next transaction are solicited. If current bid is less than current ask, then they are carried over to the next round of bids and asks for the same transaction. Since bids and asks of all traders in a given round are considered all at once, the order in which they are solicited does not matter. Since only the highest of all bids and lowest of all asks submitted in a given round matter, the order in which they are picked for comparison is not relevant. In a ***continuous*** double auction a trader is randomly solicited and his bid/ask is used to update the current bid/current ask. In solicitation without replacement the next trader is randomly picked from traders that have not been picked in a round. The process continues until the current bid exceeds or equals current ask. If all traders have been picked, the next round of random order solicitation of bids/asks for the transaction starts afresh. Solicitation *without replacement* ensures that every trader has a chance to submit a proposal before someone has the chance to submit a second proposal. Alternatively, random sampling in a *continuous* double auction could be carried out *with replacement*. Such sampling still gives equal chance to all traders, but allows for the possibility that one trader could be solicited many times while others may not be solicited at all. Under this sampling scheme, all rounds of a transaction merge into a continuum until the transaction is completed. A third possibility is to create a ***semi-continuous or hybrid*** mechanism by combining the batched bid/ask solicitation feature of *synchronized* double auction with sequential matching feature of *continuous* double auction. Similar to a synchronous auction first all bids/asks are gathered at the clearing house. Then each of them is randomly selected without replacement for updating the current bid/current asks. If the current bid at any stage of this process matches or exceeds the current ask, a transaction is completed, the remaining bids/asks are discarded, and the auction moves to the next transaction; if no match occurs, the auction moves to the next round.

4.2.3 Computerized Double Auction Market Rules

Double auctions can be run under a variety of rules. The simulation of a CDA here utilises a variation of the "semi-continuous double auction without replacement" which is described previously in Section 4.2.2. Previous researchers have mostly modelled the CDA as a synchronized Double Auction with other simplifying rules including, each bid/ask and transaction being valid for a single unit. Most previous DA studies also endow traders with the capacity to trade multiple units. This has the advantage of obtaining relatively larger data sets in terms of the number of trades taking place. In this experimentation however, partly as a result of the proposed comparative study with the Quote Driven market and also to disregard any auxiliary assumptions (such as that multiple unit traders construct strategies separately for every unit [Cason and Friedman 1995]) the experimentation employs traders with the capacity to buy/sell only single unit per trading period. An important consequence of single unit trading in our sessions with approximately 20 traders is that the markets are relatively 'thin' with few trades taking place.

A transaction cancels any outstanding bids/asks. The bid/ask prices converge until one of the parties accepts an offer from the opposite side. We also incorporate the NYSE improvement rule, which states that any new bid/offer must better the existing one in the market and in doing so it replaces the same.

4.2.4 A Pseudo code description for the CDA Protocol Simulation

0) session_no=1, round_no; r = 0

1) A new round of the CDA starts, r=r+1, ***lowAsk*** = ∞ (In the experiment this is set to 999), ***highBid*** = 0.

2) Several situations might arise in a round

   a. When a seller agent submits an ask ***a*** (i.e. the minimum price at which it will sell the item);
   -if ***a*** > = ***lowAsk*** then ***a*** is an invalid ask.
   -if ***highBid*** < ***a*** < ***lowAsk*** then ***lowAsk*** is updated to **a**;
   -if ***a*** <= ***highBid*** then the seller makes a deal at ***highBid***; GOTO 1

   b. When a buyer agent submits a bid ***b*** (i.e. the maximum price at which it will buy an item);
   -if ***b*** < = ***highBid*** then ***b*** is an invalid bid.
   -if ***highBid*** < ***b*** < ***lowAsk*** then ***highBid*** is updated to ***b***;
   -if ***b*** >= ***highBid*** then the buyer makes a deal at ***lowAsk***; GOTO 1

3) Step 2 repeats for a certain number of rounds (till a session ends). In each round all buyers/sellers submit a bid (ask) once only.

4) The experiment is repeated with session_no = session_no + 1 and new schedules for buyers/sellers and reinitialised initial conditions except the learning parameters. One experiment is comprised of a fixed number of market sessions.

The outstanding lowAsk and highBid [highBid, lowAsk] define the bid/ask spread and only bids and asks that fall within this region are considered valid.

4.2.5 The Adaptation Algorithm

The price adaptation algorithm utilised by our traders is outlined here. It is a variation of that proposed in the works of Cliff and Preist [Cliff 1997; Priest and Tol

1998] and outlines the procedure for price adaptation in our DA market set-up by the buyers and sellers.

The ZIP traders operate by maintaining a ‘profit margin’, which determines the price they quote (bid or offer) in the market. The objective of each ZIP trader is to maximise the gains from trading i.e. their profit. This involves deciding on:

a) When to increase (or decrease) the profit margin and

b) How much to increase (or decrease) the profit margin by.

Following on from Cliff and Preist [Cliff 1997; Priest and Tol 1998], a trader makes the qualitative decision on when to alter its profit margin based on the factors below:

i. If the trader is active i.e. whether it has bought or sold its entitlement. A trader becomes inactive after it has bought or sold its quota; however it is still allowed to update its profit margins.

ii. If the last quote by a trader was a bid or offer.

iii. If the bid or offer was accepted i.e. if it lead to a trade.

iv. If the last quoted price was greater than or less than the price the trader in question would currently quote.

The quantitative question of how much to increase (or decrease) the profit margin by is described next.

Each trader is initially allocated a profit margin (μ) and the price (p) it quotes at each time step is based on its limit price (λ) and its profit margin at that instance, i.e.

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

At each time step, a trader determines a target value towards which it attempts to adjust its price (by altering its profit margin). This is achieved as below.

Let $\mathbf{B_{max}}$ be the highest bid in the current round, prior to trades taking place, and $\mathbf{S_{min}}$ be the lowest offer. Let d be a random value, small with respect to $\mathbf{B_{max}}$ and $\mathbf{S_{min}}$. The target value t for agents to adjust towards is determined as follows:

For BUYERS:

If $\mathbf{S_{min}} > \mathbf{B_{max}}$ then

target = $\mathbf{B_{max}} + d$

If $\mathbf{S_{min}} <= \mathbf{B_{max}}$ then

target = $\mathbf{S_{min}} - d$

For SELLERS:

If $\mathbf{S_{min}} > \mathbf{B_{max}}$ then

target = $\mathbf{S_{min}}$ - d

If $\mathbf{S_{min}} <= \mathbf{B_{max}}$ then

target = $\mathbf{B_{max}}$ + d

For our experiments, we follow [Cliff 1997] in our definition of d:

If the target is $\mathbf{B_{max}}$ + d

then d = r1 $\mathbf{B_{max}}$ +r2

If the target is $\mathbf{S_{min}}$ - d

then d = r1 $\mathbf{S_{min}}$ +r2

where r1 and r2 are independent random variables identically distributed in the range [0,0.2].

Given the target value, the agent does not jump straight to that value, but moves towards it at a rate determined by the learning rule. The learning rule used is Widrow-Hoff with momentum [Cliff 1997; Priest and Tol 1998], which is also used for back propagation learning in neural networks. The learning rule has two parameters. The learning rate β determines the speed with which the adjustment takes place, and the momentum У acts to damp oscillation. Given p(t) and t(t), the valuation and target price at time t, the learning rule determines a new profit margin , μ (t+1), as follows:

$$\mu (t+1) = (p(t) + \Gamma(t))/ \lambda - 1$$

where $\Gamma(t + 1) = У\Gamma(t) + (1 - У)ß(t(t) - p(t))$ and $\Gamma(0) = 0$ [Cliff 1997]

4.3 The Simulation Infrastructure

A custom simulator written by the author is used for all the experiments. The Simulator is developed in Java and runs on a single PC. The code is organised into 3 main classes and a variety of auxiliary functions.

1) Agent.java – The Buyers/Sellers in the market place extend the agent class.
2) Market.java – In the simulation the Market polls Buyers/Sellers at each time step. The buyers/sellers run their own internal cycle and are in knowledge of the current best bid/ask in the market. On being polled if an agent is 'active'

it runs its internal update cycle and sends a bid or ask to the market. The market only accepts bids/offers that improve upon current ones.

3) Simulator.java – This class runs the simulation tool and when called upon, starts a new instance of the Market. The number of buyers/sellers and their activity rates can be specified through a GUI.

4) Besides the above classes, there are other supporting classes which define the messaging structure, statistical distributions used etc.

The results presented here are averaged over 50 experiments[25] with 100 market sessions in each experiment. The data is interpreted via Matlab scripts and the graphs are plotted using Matlab and Excel.

The figure below represents the demand supply conditions existing in the marketplace which 11 buyers and 11 sellers. Theoretical analysis predicts the competitive price as 4.75 units and the competitive quantity traded to be between 5 to 6 units. The total surplus (profit) available to the market participants is determined as the area to the left of the demand supply intersection and is 7.5 units as shown. This surplus should ideally be distributed equally among buyers and sellers in a perfectly competitive market[26].

Fig. 4.1 Demand/Supply conditions for the market simulation

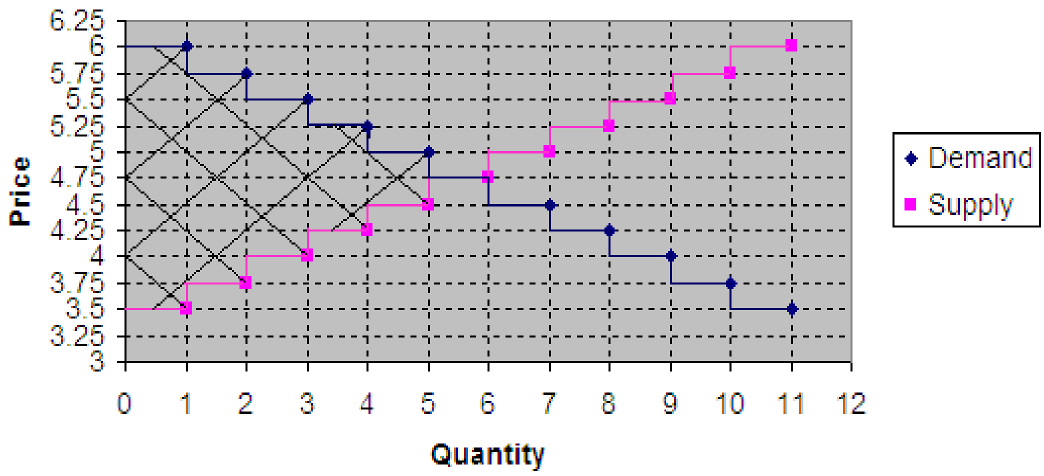

---

[25] This value is chosen because it provides statistically significant results. For example, trader profits from 50 and 100 experiments differ by only 0.004, which is not significant in this case.

[26] The fixed number of units up for trade and the fixed reservation prices dictate a stepped demand-supply curve for the market outlined.

4.4 CDA: Results

The figures below shows the transaction price series for one experiment of the CDA market with ZI-C and ZI-U traders. The ZI-C traders make random bids/offers restricted only in that they will not make a bid/offer at a loss i.e. they don't implement any learning mechanism over time, but do take care to bid randomly within their limits. The ZI-U traders on the other hand bid in a completely random fashion. A plot of ZI-C and ZI-U prices shows that prices do not converge across sessions, but within a session, prices do converge in a ZI-C market[27]. Over 50 experiments in a ZI-C setting, on average the last transaction price in each session is lower than the first one by about 7%. However there is no such inter session price convergence in a ZI-U environment. These results again validate those obtained by Gode & Sunder [Gode and Sunder 1993] and later by Cliff [Cliff 1997].

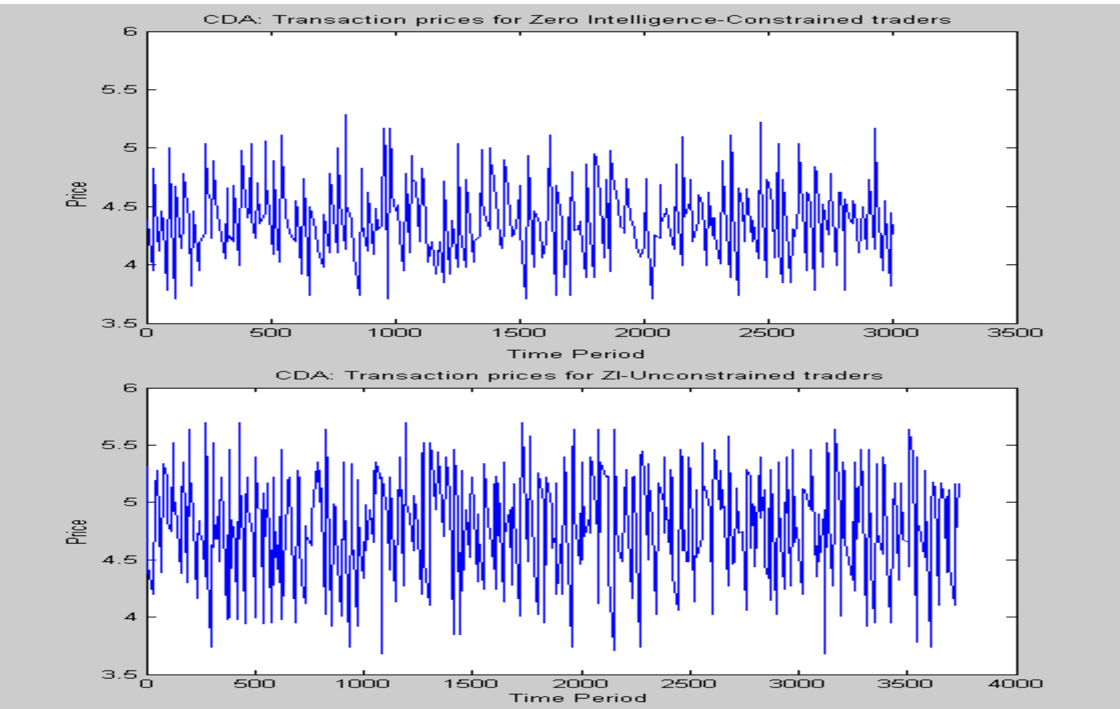

Fig. 4.2 Top: Trade prices in a CDA with ZI-C traders Bottom: with ZI-U traders

27 Gode & Sunder explain this convergence as a result of the progressive narrowing of the feasible range of transaction prices as more units are traded during a session.

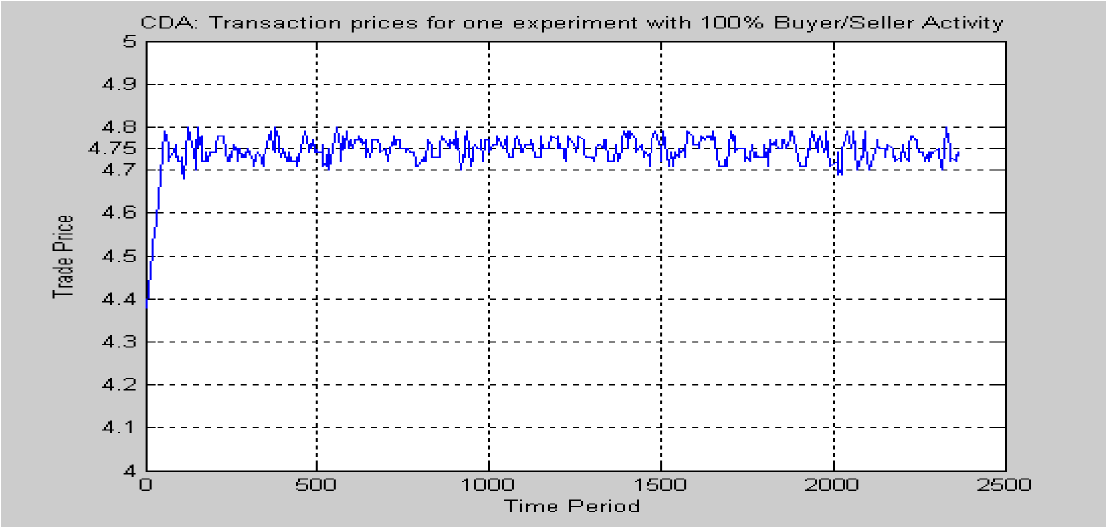

Fig. 4.3 Trade prices in a CDA with ZIP traders which follow the rudimentary learning strategy outlined previously

In contrast to a transaction prices in a CDA market seen with ZI-C and ZI-U traders, a plot of the same with the ZIP traders implementing a learning strategy shown above clearly indicates that prices converge very rapidly to equilibrium (4.75 units) and stay close to it during the period of the experiment. Within a session too, there is still further convergence as seen in the ZI-C environment. The figure below plots the transaction prices within a session as a result of the bids and offers converging.

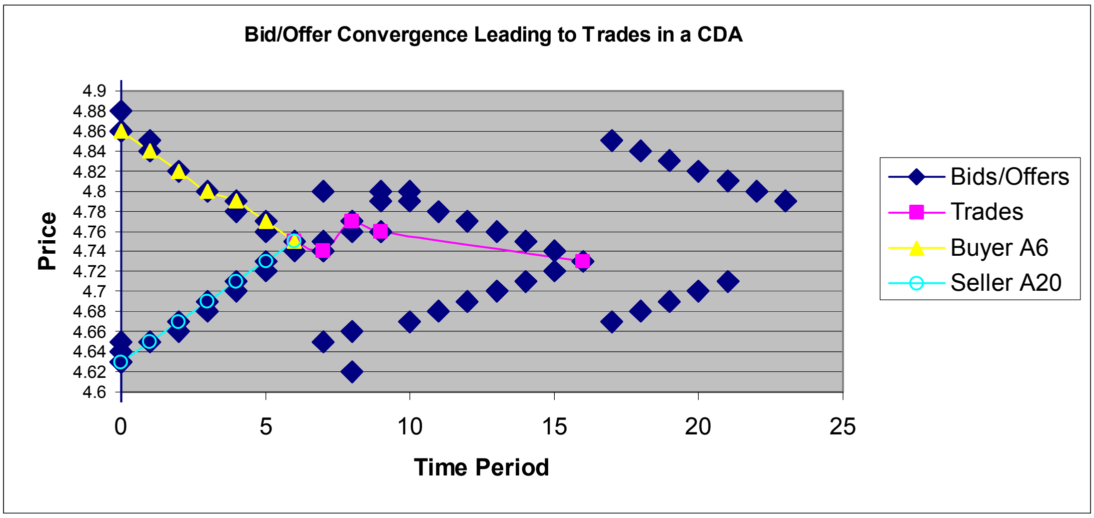

Fig. 4.4 Trade prices within a single session with bids and offers converging.

Only a few accepted bids/offers (that improve upon previous bids/offers in the market) are shown for clarity. The figure is based on synchronized bidding with all buyers/sellers able to bid at every time step. A chronology of the first trade between Buyer A6 and seller A20 is shown. Since both the buyers and sellers are following the same learning strategy, the convergence of prices is almost symmetric close to equilibrium levels.

The results that follow show the variation of α (Alpha)[28] with session number averaged over 50 experiments as also the profit loss distribution for the traders and overall market efficiency as determined from the total surplus extracted.

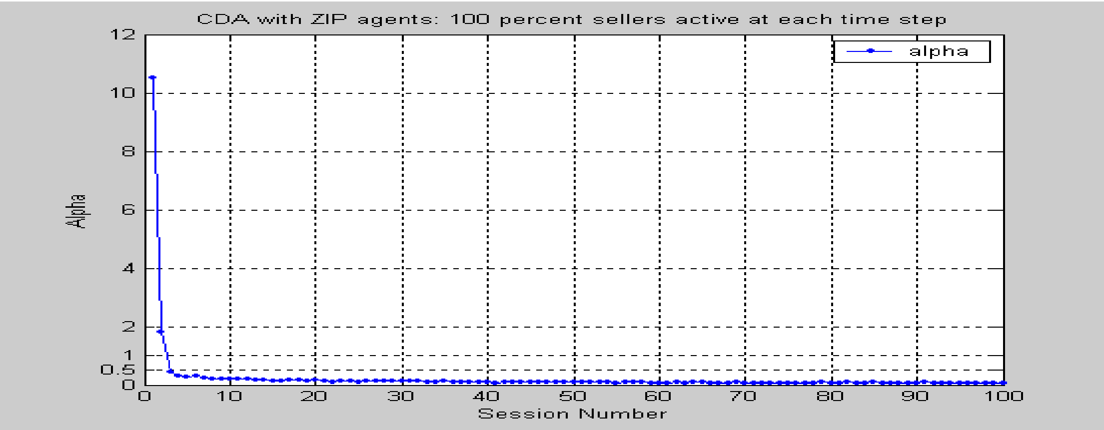

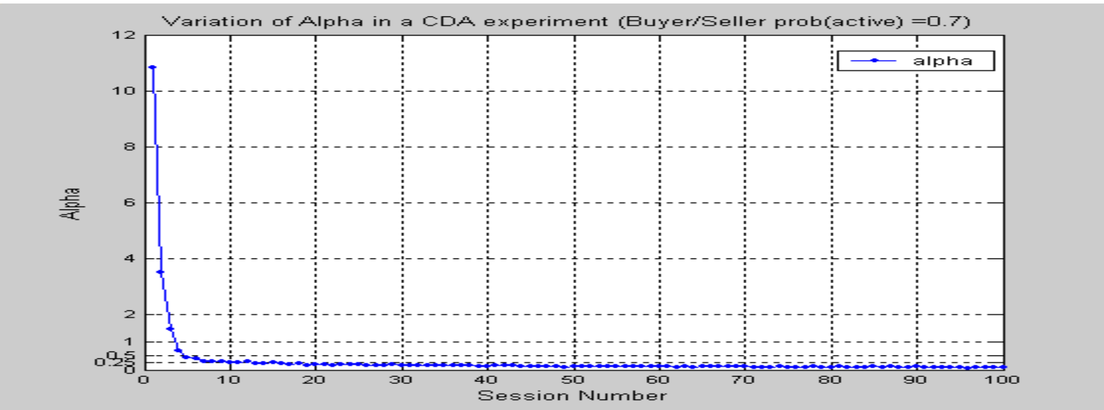

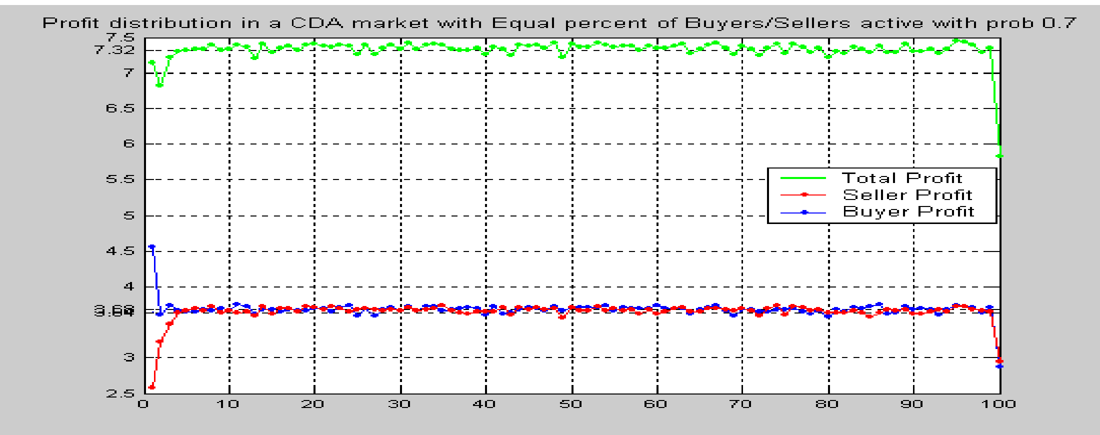

[28] For definition of Alpha see page 38.

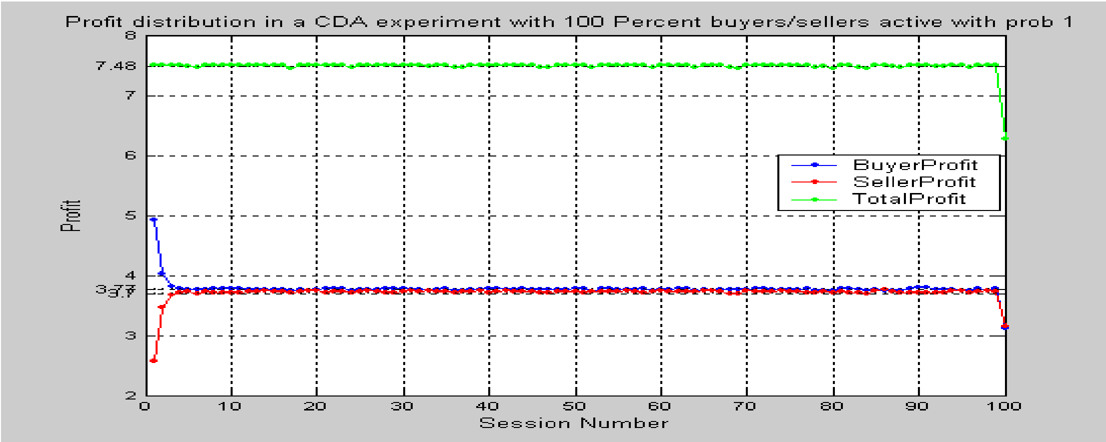

Fig. 4.5 Variation of Alpha and Profit for Buyers/Sellers in a market with 100% buyer/seller activity and one with 70% activity for both buyers and sellers

The figure on the previous page validates our model of the CDA with respect to two significant attributes outlined previously in the work of other researchers. These are the convergence to equilibrium and observed efficiency of markets organised with the Double Auction protocol. The CDA has for long been known to encourage quick and reliable convergence to competitive equilibrium in a wide variety of laboratory environments [Cason and Friedman 1995]. The results on the previous page can be compared against the work of Priest et al [Priest and Tol 1998] who model a CDA as a synchronized bidding environment (i.e. all the agents are allowed to make a bid/offer at each time step. This is captured in the figures with 100% buyer/seller activity). The author has also compared a slightly different environment proposed by Tesauro & Das et al [Das, Hanson et al. 2001; Tesauro and Das 2001]. In their work, each agent (buyer/seller) has a constant activation probability per time step different from 1. The author has modelled a similar situation with experiments giving the buyers/sellers with varying constant probabilities of activation. The results presented on the previous page are for a constant probability of 0.7. From the figures it is clear that the results are qualitatively similar as long as the probability of activation is same for both buyers and sellers. The only difference is that in a market with probability of activation less than 1(i.e. 0.7), the overall surplus extracted (and hence market efficiency) comes down to 7.32, from 7.48 observed in a market with probability 1. This drop can be attributed to some profitable trades not taking place because only a certain (and random) population of traders is active and allowed to trade at each time step. The figures compare with a theoretical value of 7.5 that is predicted with the demand supply conditions existing in the market. Hence the CDA market provides close to 100% efficiency. The available surplus (profit) is also distributed equally among buyers and sellers, in line with theoretical predictions. A plot of Alpha, which provides a measure of the convergence to equilibrium, is also in line with previous results (see from Priest et al). Prices converge rapidly to equilibrium and Alpha is below 1 within the first few sessions.

The next few experiments are also plots of Alpha vs. session number and the profit distribution for the traders. The difference from the approach of previous researchers is that the author has tried to vary demand/supply[29] in the market by providing the sellers with varying probabilities of activation keeping the buyer activation probability constant at 1 (and vice versa). Through this approach we aim to observe the performance of a CDA market and attempt to capture the salient features, which can be used to compare the CDA with a Quote Driven model presented later.

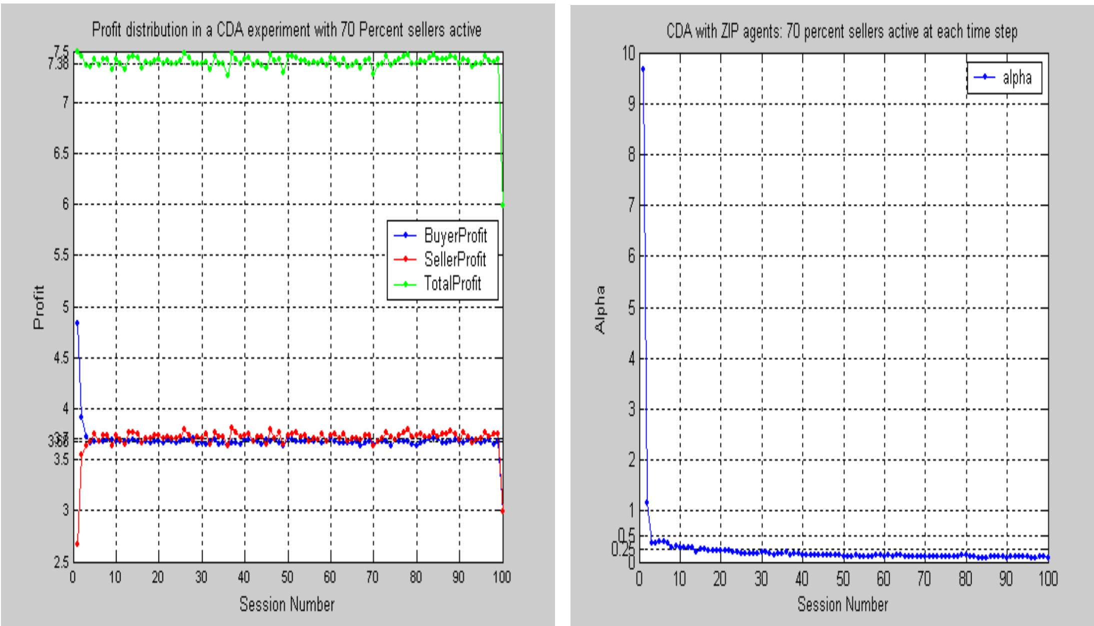

Fig. 4.6 Left: Plot of profit distribution in a CDA market with only 70% sellers active (all buyers are allowed to bid) Right: Plot showing variation of Alpha averaged over 50 experiments. It is instructive to see that the profits remain fairly equally distributed. However the total profit extracted comes down to a mean of 7.38 units. Alpha also converges rapidly to below 1 within a few time steps.

29 Previous researchers vary demand and supply by changing the limit prices of the commodities available to agents during the course of an experiment. This only serves to change the equilibrium price and in some cases the quantity traded. However the purpose behind this approach is simply to show that the adaptation algorithm identifies the underlying change and is able to track it. It does not aim to determine the affect of such a change on observables such as market efficiency.

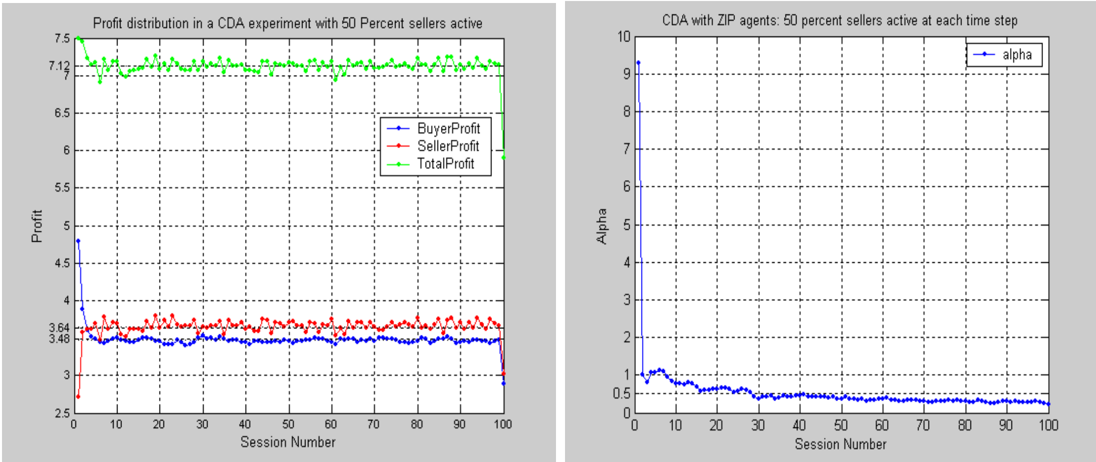

Fig. 4.8 Left: Plot of profit distribution in a CDA market with only 50% sellers active (all buyers are allowed to bid) Right: Plot showing variation of Alpha averaged over 50 experiments. Although buyer profits have fallen, efficiency is still > 90%

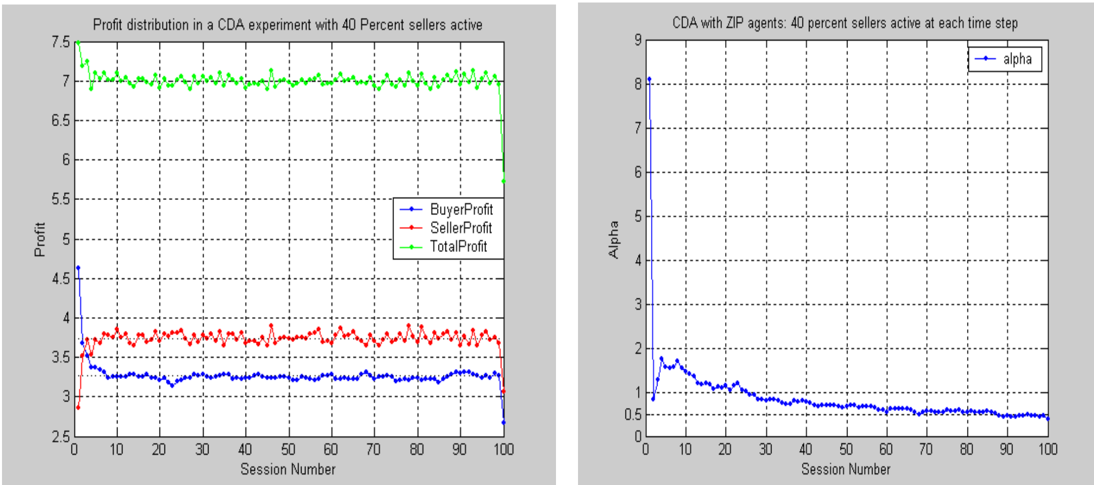

Fig. 4.9 Left: Plot of profit distribution in a CDA market with only 40% sellers active (all buyers are allowed to bid) Right: Plot showing variation of Alpha averaged over 50 experiments. Efficiency ~ 93% but Alpha takes longer to settle below 1.

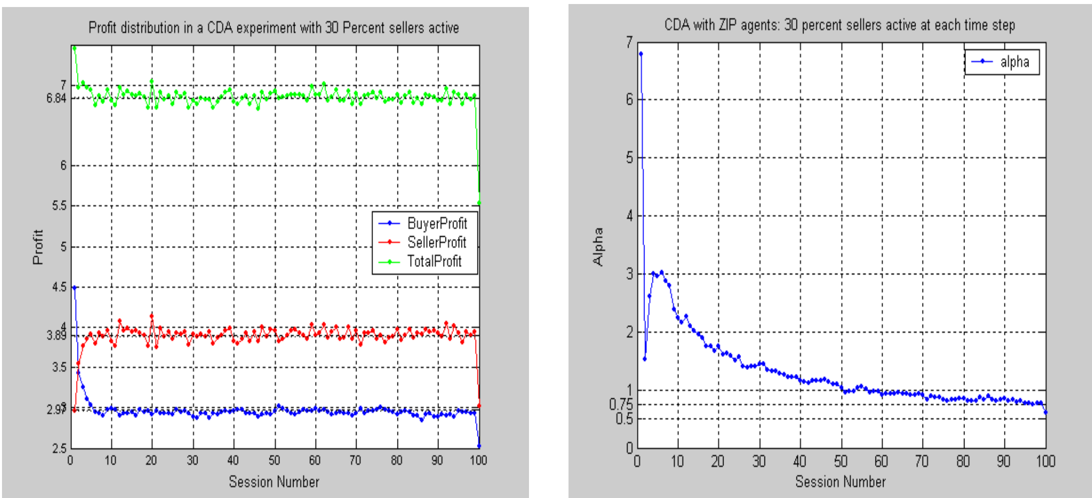

Fig. 4.10 Left: Plot of profit distribution in a CDA market with only 30% sellers active (all buyers are allowed to bid) Right: Plot showing variation of Alpha averaged over 50 experiments. Profits for sellers increase but overall efficiency falls and it takes almost double the number of sessions for Alpha to fall below 1 than with 40% seller activity

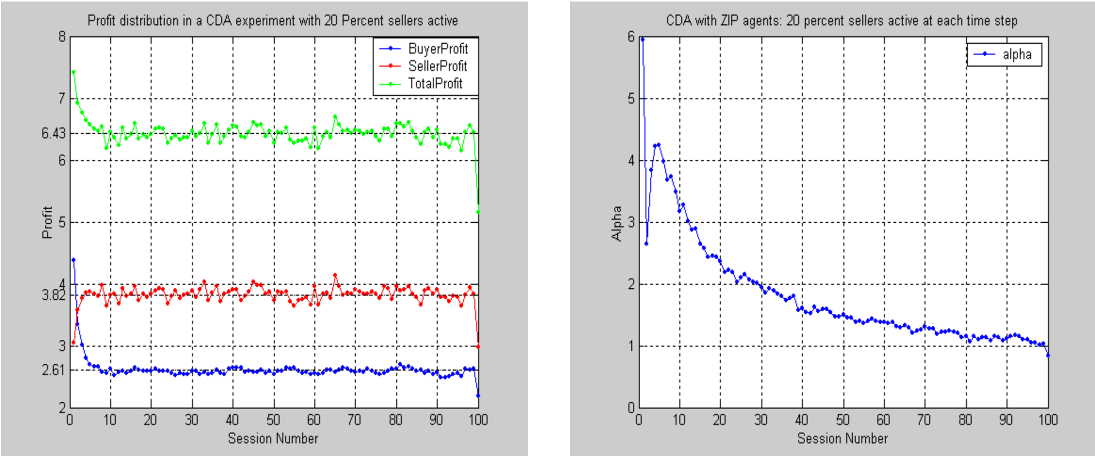

Fig. 4.11 Left: Profit distribution in a CDA market with only 20% sellers active (all buyers are allowed to bid) Right: Variation of Alpha averaged over 50 experiments. Alpha now takes ~ 100 sessions to fall below 1.

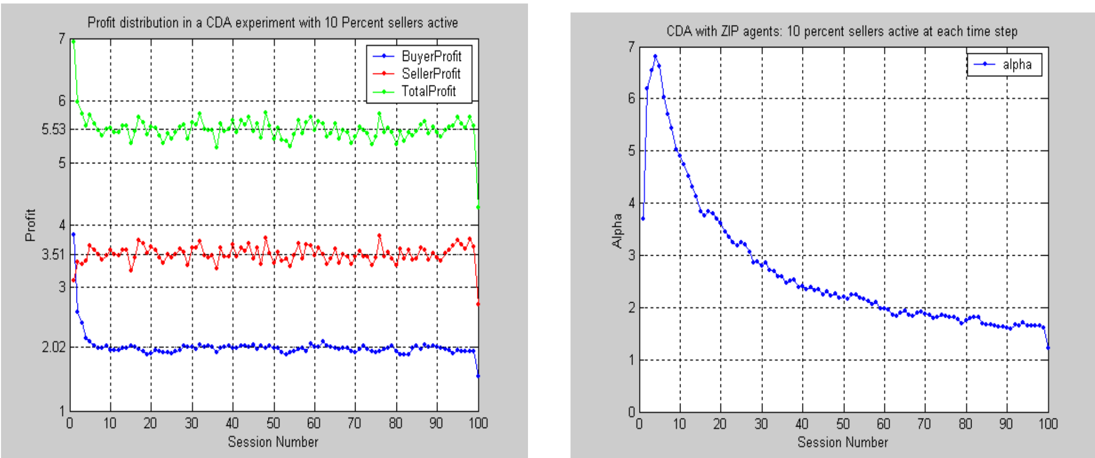

Fig. 4.12 Left: Profit distribution in a CDA market with only 10% sellers active (all buyers are allowed to bid) Right: Variation of Alpha averaged over 50 experiments. Note that Efficiency falls to 74% and Alpha does not converge to values below 1 within 100 sessions. Also note that seller profits are falling too with buyers

The following pages show plots of repeat experiments with buyer activity rates being varied while keeping the seller activity rate constant at 1. The plots are qualitatively similar to those obtained previously for varying seller activity.

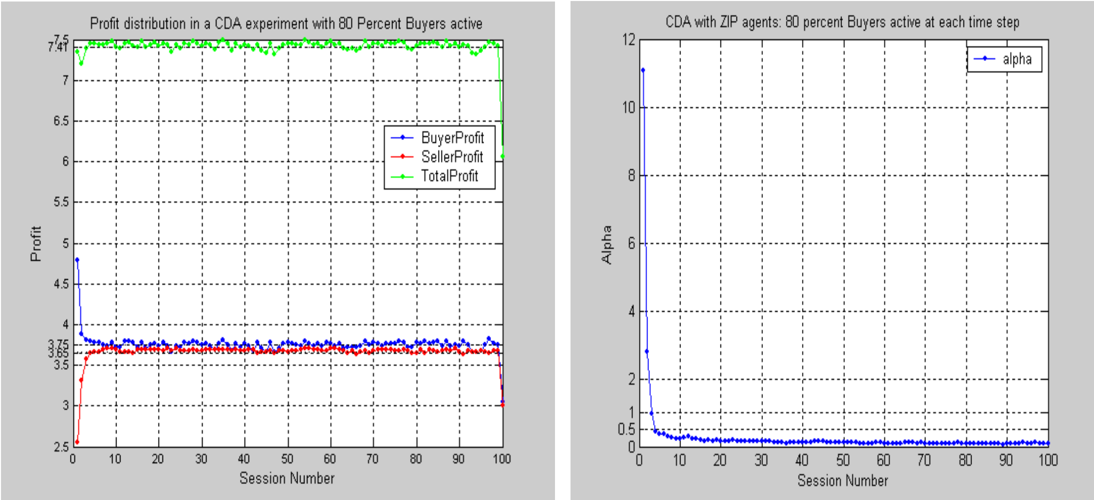

Fig. 4.13 Left: Plot of profit distribution in a CDA market with only 80% buyers active (all sellers are allowed to bid) Right: Plot showing variation of Alpha averaged over 50 experiments.

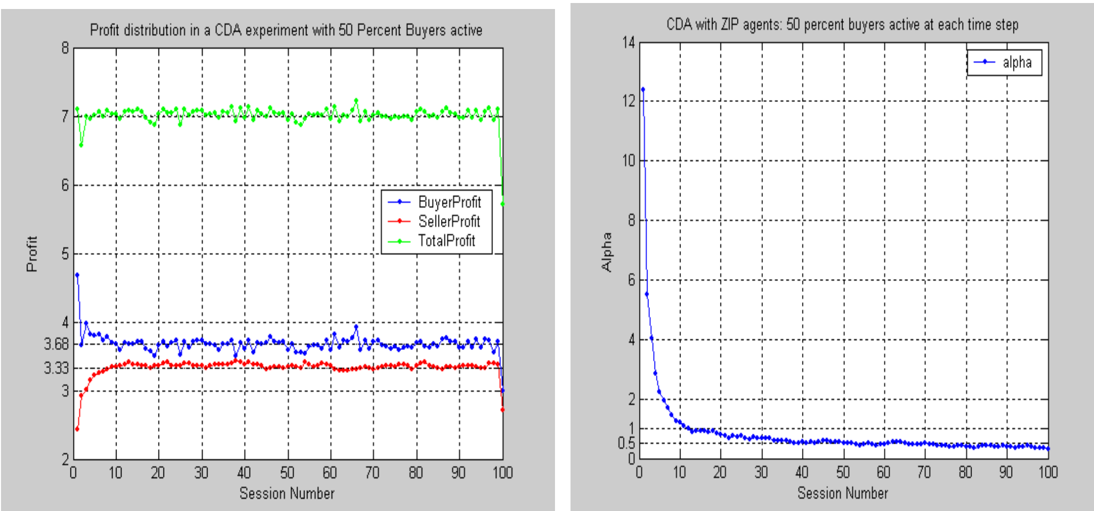

Fig. 4.14 Left: Plot of profit distribution in a CDA market with only 50% buyers active (all sellers are allowed to bid) Right: Plot showing variation of Alpha averaged over 50 experiments. Note that seller profits are falling in these experiments.

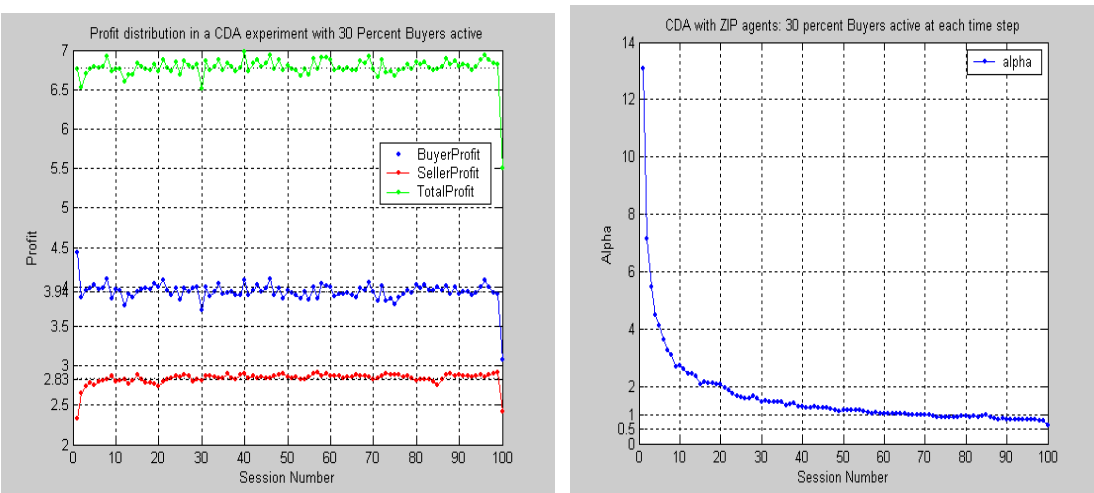

Fig. 4.15 Left: Plot of profit distribution in a CDA market with only 30% buyers active (all sellers are allowed to bid) Right: Plot showing variation of Alpha averaged over 50 experiments.

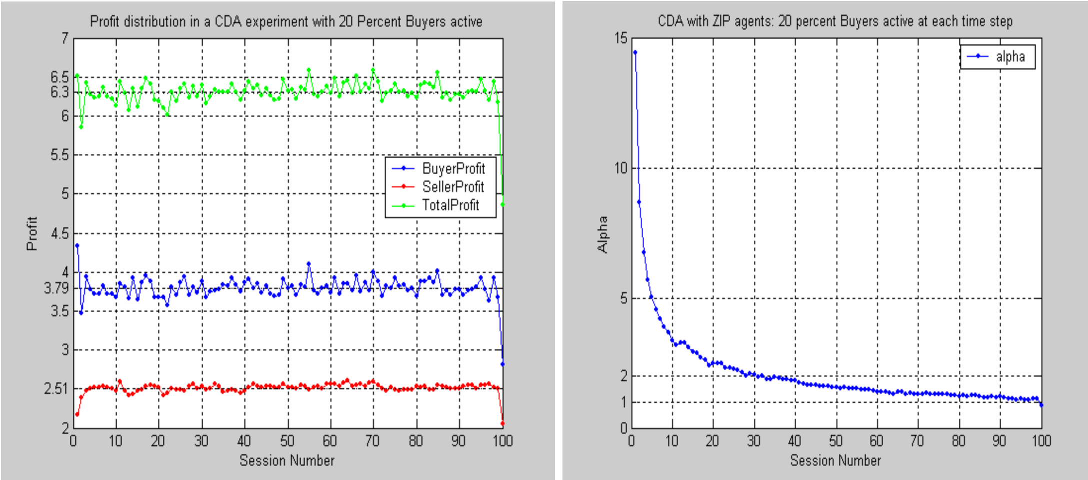

Fig. 4.16 Left: Plot of profit distribution in a CDA market with only 20% buyers active (all sellers are allowed to bid) Right: Plot showing variation of Alpha averaged over 50 experiments.

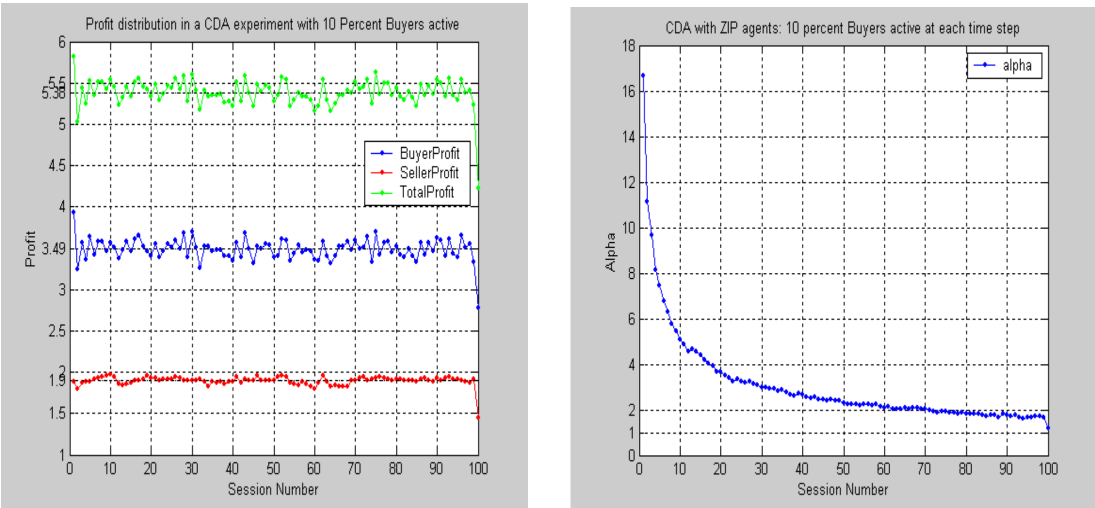

Fig. 4.17 Left: Plot of profit distribution in a CDA market with only 10% buyers active (all sellers are allowed to bid) Right: Plot showing variation of Alpha averaged over 50 experiments. Note that Alpha does not fall below 1 for the 100 sessions.

An analysis of the previous results indicates results consistent with theoretical predictions. It also reveals that the CDA market is quite robust to changes in demand supply conditions. Even a 50% change in buyer (seller) population still leads to prices near to those predicted in a competitive equilibrium (captured in figure 4.18). The figure below shows median values of transaction prices obtained in a CDA market with varying seller activity.

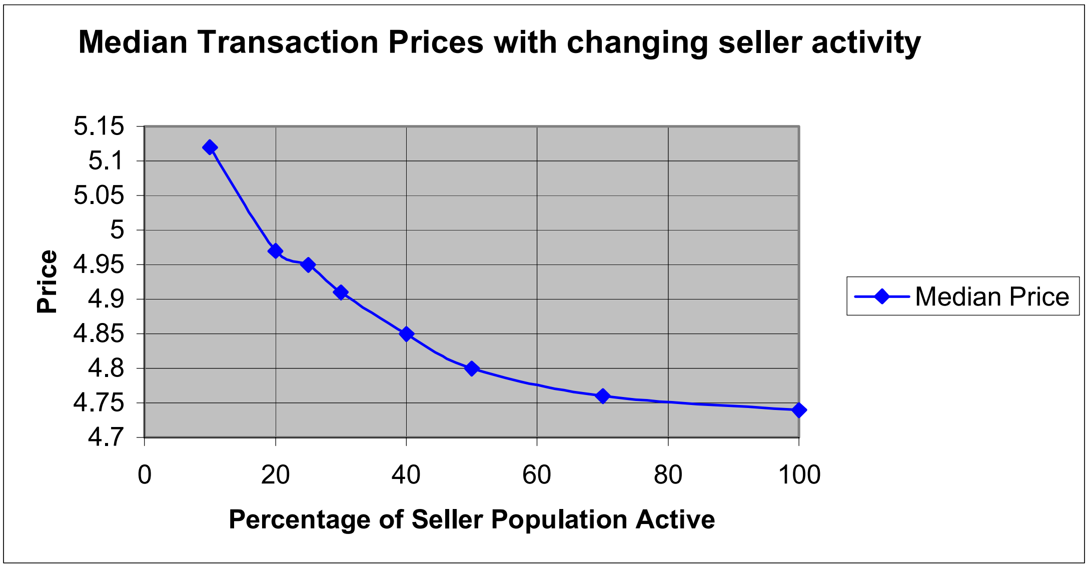

Fig. 4.18 Median value of transaction prices with changing seller population.

However buyer/seller activity below 30% leads to dramatic falls in both the equilibrium prices obtained and the efficiency (in terms of surplus extracted) achieved from the market. We next try to verify the reasons behind this phenomenon.

The figures below plot the variance in buyer and seller profits for the sets of experiments conducted previously. It is evident that even with buyer seller activity reduced to 50% (compared to the other), the surplus in the market is almost equally distributed i.e. neither of the buyers or sellers are particularly disadvantaged. However below 50% buyer (seller in the following plot) profits fall down and seller (buyer) profits tend to increase. This is consistent with the theory of supply and demand. Fewer sellers in the market will obtain a better price for their goods than a situation where price competition forces their prices down and similarly for the case of buyers in the next plot.

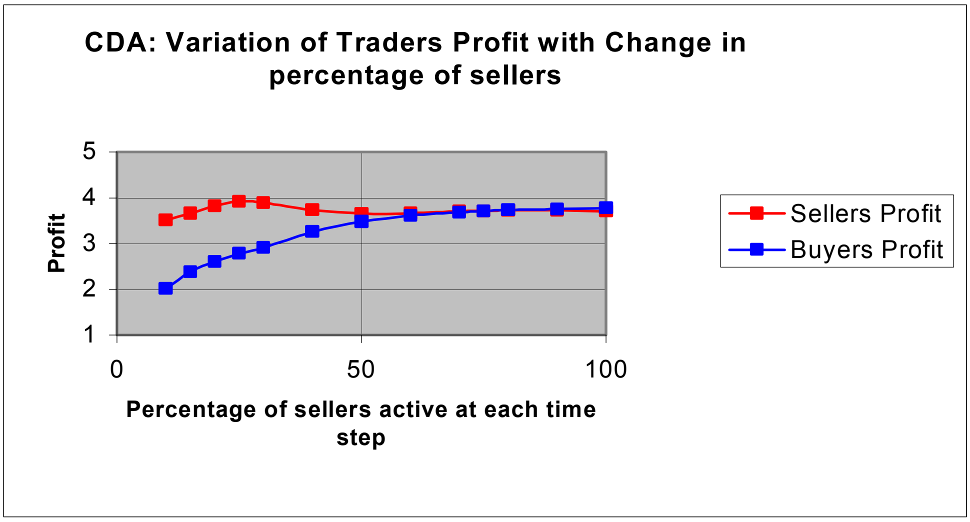

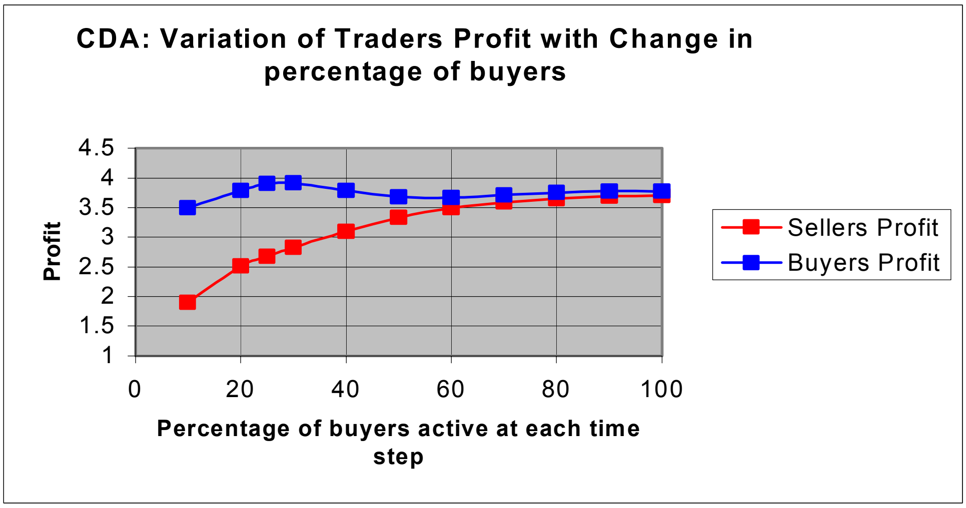

Fig. 4.19 Top: Variation of trader's profit with changing seller activity Below: Same experiment repeated with changing buyer

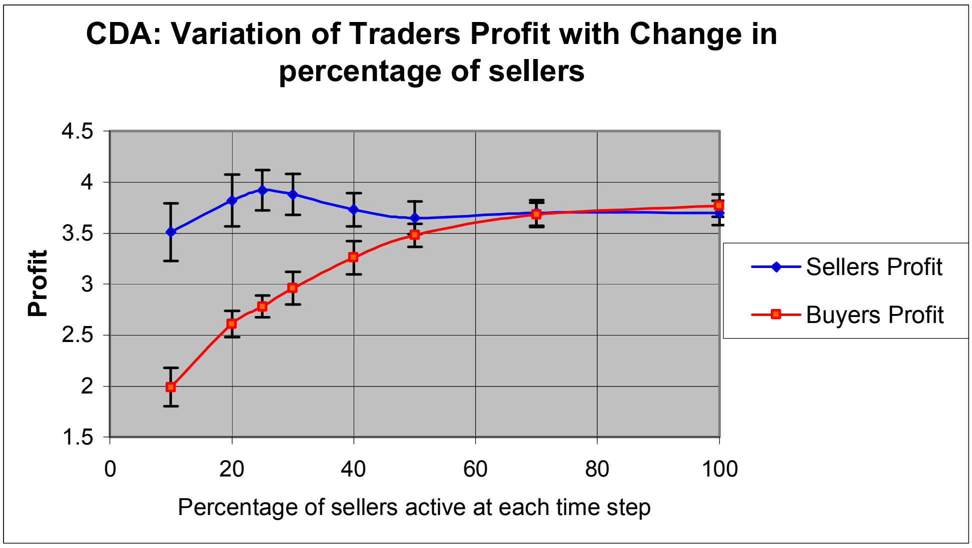

Fig. 4.20 Top: Error bars plotted as the 95% confidence interval for the buyer/seller profit margins.

One peculiarity observed in the above graphs however is that both buyer/seller profits tend to fall, below a 25% activity for the opposing side. While buyer profits (seller in the next case) continue to fall as expected, the fall of seller (buyers in the next) prices is remarkable. A reason for this occurrence can be seen in the figure below which plots on average the total number of trades occurring in the CDA market when the seller activity rates are varied.

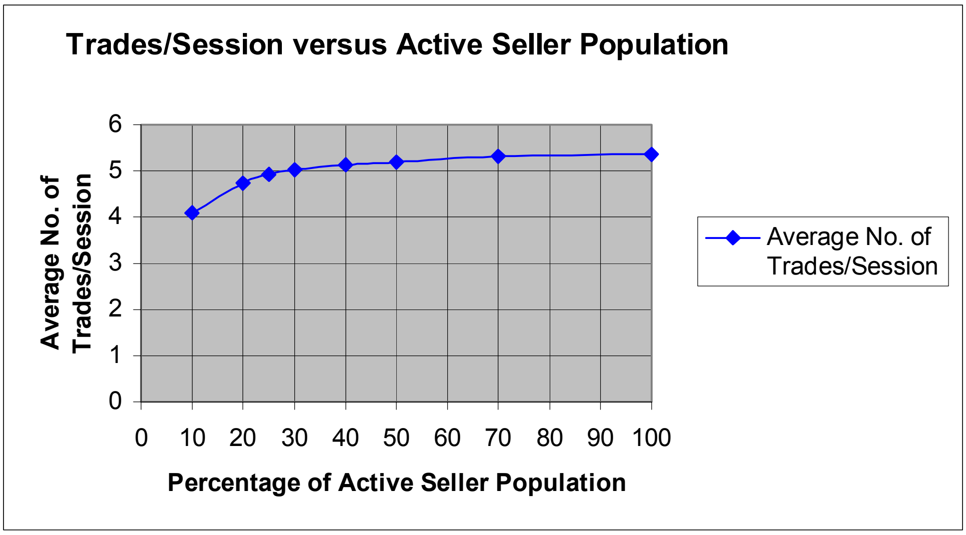

Fig. 4.21 Trades per session averaged over 50 experiments for varying seller activity.

Remember that from the demand supply conditions attributable in the market due to our experimental setup, there should be on average between 5 to 6 trades per session. From the plot below it is clear that above 30% seller activity, this condition is maintained. However activity rates below 30% lead to the market seeing on average less than the 5 trades per session predicted by competitive market theory. The reason for this can be explained in the way that our experiments are set up. We initially proposed to limit the length of a session to 3 consecutive rounds of non-improvement in the bid/offer prices. However with a drastically reduced buyer (seller) population, the price adaptation process becomes slower and as such with sessions ending within 3 rounds of non-improvement in prices, a small number of potentially viable trades are not conducted. Further experimentation (not shown here) with increasing the number of rounds with non-improving prices to 20 (to mark the end of a session) moves the number of transactions to between the competitive range of 5-6 again.

Figure 4.21 captures further interesting characteristics from the CDA market scenarios implemented above. Here we consider the graph of trade prices observed during one experimental run of the CDA market with 50% seller activity (100% buyers). It is interesting to note that although the average transaction prices stay close to the competitive equilibrium i.e. 4.75 (for the fig with 50% seller activity at least), there is increased volatility in the prices (i.e. the standard deviation of prices is far greater than a run of the CDA market with 100% buyer/seller population as in **Fig.** 4.3). This result is important in the sense that one of the reasons for implementing Quote Driven markets in the real world (and indicated previously as an advantage in this report) is to reduce price volatility (besides offering a continuous trading environment). This is because prices are determined solely by customer (trader) orders in a Double Auction market. In a Quote Driven market maker, the middleman has certain obligations to provide a stabilising effect from huge price swings.

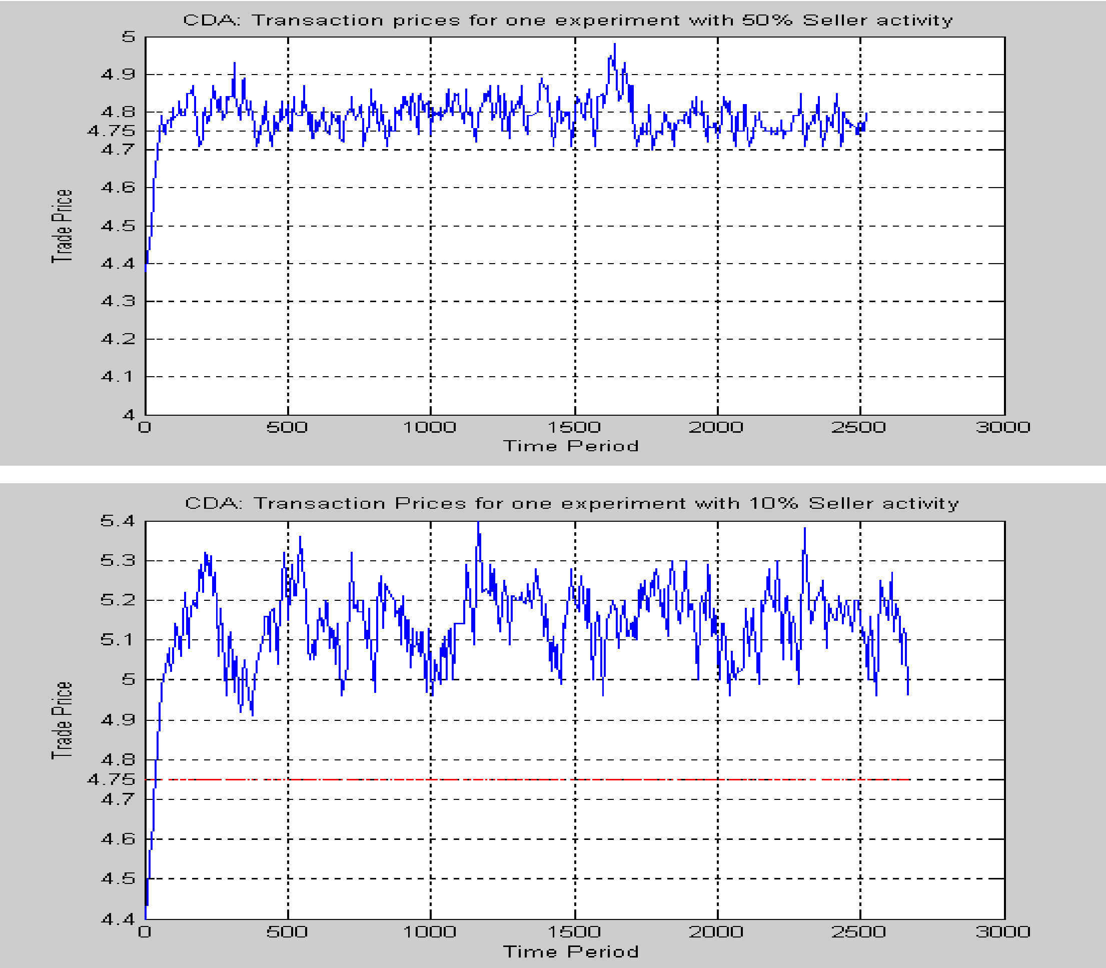

Fig. 4.22 Plots of trade prices in 2 experiments with 50% and 10% seller activity. Note the fluctuation of the prices around the equilibrium price when compared to fig 4.3

4.5 Quote Driven Market: An Experimental Approach

This section gives an outline of an experimental approach that could be followed to compare the CDA and Quote Driven markets.

The Quote Driven market would also incorporate the same hierarchy of *market*, *sessions*, *rounds* of bidding and *transactions* to enable a comparative assessment with the DA market. The market organisation will be different with all buyers/sellers having to transact against the bid/offer price indicated by the market maker. We attempt to construct the simplest possible Quote Driven scenario, with a single, monopolistic, risk neutral market maker who sets prices, receives all orders and clears trades. The market makers objective is to maximise expected profit per unit time. The buyers/sellers in the marketplace employ the same ZIP learning strategy used in the CDA market implementation.

The market variables which the market maker takes into account for adjusting prices are dependant on:

1) The inventory of the market maker
2) Order Imbalance: i.e. difference between the buy/sell orders received within a period of time

The options, which it has with respect to its actions, include:

1) Change the bid price
2) Change the ask price
3) Set the bid size
4) Set the ask size (bid and ask sizes can be fixed at 1 unit for comparison with a CDA market as outlined above)
5) Others (Buy or sell, provide price improvements i.e. better prices than current market quotes)

What follows here is a description of a simple Quote Driven market set-up which can be utilised in a comparative study against the CDA market. We can simplify the available options (available to the market maker as above) by stipulating a fixed bid-ask spread. This means that the market maker can essentially change his bid/ask prices keeping the spread same. The first session in the market place is similar to a

DA market i.e. all buyers and sellers are expected to make a bid/offer. The market maker initially sets his prices based on the prices obtained in this 'opening auction'[30]. These prices are then varied in accordance with the arrival of buy/sell orders in later periods. An important variable in the design of the Quote Driven market is whether the price quotes/trade prices are transparent (i.e. available to all market participants) or there can be varying degree of opaqueness (i.e. the trade prices may not be available to participants or the bid-ask prices can only be only be obtained through bilateral quoting (i.e. the interested buyer/seller will need to contact the market maker to obtain the price). Initially we presume that a simple implementation of the market could be fully transparent i.e. both the bid-ask prices and the trade prices are publicly available.

The market maker is assumed to operate with simple rules of thumb; following a near zero drift inventory policy i.e. if the previous transaction is a buy then the market maker will adjust it's prices in the next period to obtain better prices for a buy or conversely to increase the chances of a sell. As a consequence, the optimal bid/ask prices posted by the market maker are monotone decreasing functions of its inventory position. As the market maker's inventory increases, he lowers both bid and ask prices and conversely he raises both prices as inventory falls (Such a link between prices and inventory is shown in the early models of market maker behaviour e.g. [O'Hara 1995]).

---

30 Most financial markets have an opening and closing auction separated by a period of continuous trading. *Specialists* on the NYSE set their prices based on the knowledge of the prices in the opening call auction.

Chapter 5

Discussion & Conclusion

5.1 Overview

With the recent advances in e-commerce, using software agents to trade on our behalf may well be the norm rather than the exception. This requires an understanding of the markets in which the agents will operate. The markets will need to be designed so as to create an accessible and fair environment in which the humans will be comfortable with the operation of their personalized agents. By offering a comparative simulation study between two widely used continuous market mechanisms that support trading between multiple buyers and sellers, this work aims to provide an insight into the question raised at the beginning of this thesis i.e. *What type of market environment can best achieve this situation and under what conditions?*

Throughout this thesis the self-interested behaviour of agents is emphasised, which will undoubtedly have economic decision-making capabilities. Although this thesis deals exclusively with agents for trade in commoditised markets, we can foresee a role for them, which can vary from providing information services [Kephart, Hanson et al. 2000], ontology translation, match-making, network service provision and much more. The distinctive feature about them will be that these agents will charge a fee for their goods or services and will negotiate both as buyers and sellers with other agents. Thus they will have to be economically intelligent, capable of making effective decisions about pricing, purchasing or bidding.

A crucial aspect dealt with in this thesis is the role of the mechanism or trading arrangement by which the actual process of trading is carried out in the market. Several mechanisms were considered which are suitable for different settings (types and quantities of goods, types and numbers of the agents involved etc.). For purposes of relevance, the discussion is centered on many to many trading mechanisms like the double auction in greater depth. It is in this context that the Quote Driven market has been briefly presented as a viable mechanism for automated trading where the self-interest of the agents can be accounted for, and at the same time, we can possibly accomplish the goals of global efficiency and social welfare for the domain in question.

The problem of efficient mechanism design is also discussed, which is the ultimate goal of the system designer. However this is not possible without a good understanding of the game theoretic considerations of the players involved. The design of an auction mechanism for the sale of spectrum license's for wireless communications (e.g. the recently held 3G mobile auctions) is a good example of mechanism design implementation [McMillan 1994; McAfee and McMillan 1996]. That the auction succeeded in several places including the UK but failed to generate the expected outcomes in several others just goes to show the difficulties involved in mechanism design.

Another relevant issue regarding mechanism design is its usefulness in studying the Quote Driven market. Recall that the discussion of evaluation criteria for negotiation mechanisms (Chapter 2) indicated that incentive compatibility (i.e. revealing true preferences in equilibrium) is a desired quality in an efficient mechanism. However in that case it was assumed that the elicitation of preferences (as truthful bids) was to be obtained from the bidders (e.g. the buyers in an auction). In the case of the Quote Driven mechanism, the buyers or sellers do not reveal any information to each other directly i.e. the market maker makes the first move by offering a bid-ask price. In such circumstances, the problem of mechanism design is more relevant from the viewpoint of the behaviour of the market makers. The efficiency of the market is dependant to a great extent on the competition between the market makers to achieve narrower bid-ask spreads. In the Quote Driven bandwidth market [Bourne 2000], there is bilateral communication between the SPs to determine the quoted prices. This approach could be compared against a market structure in which the quotes are publicly available price queues (see [Flood, Huisman et al. 1999]). It is in this context too that the information-based models of dealer behaviour presented in previous work [Zaidi 2001] assume significance. Given the dynamic nature of the order flow (arrival of customer orders) in the market and the market makers job of setting the quotes and executing orders based on this observation (with profit maximization as a long term objective), we can view this as a problem of *learning* or *adaptation* for the market maker. Chan and Shelton [Chan and Shelton 2001] describe a reinforcement learning model of an market maker that takes into account the inventory, order imbalance (demand supply imbalance) etc. as the environment states to determine the appropriate response for the market maker.

5.2 Experimental Analysis

5.2.1 Stressing the Novel Aspects

The development of a proprietary discrete time simulation tool has provided us with a controlled environment allowing evaluation of an otherwise complex structure as a market. The author's approach in utilising minimally adaptive agents is useful in that it provides a basis to differentiate observed results in any comparative assessment along the lines of market organization and protocols.

This thesis has presented useful and concrete data from the experimental tool, which validates the Continuous Double Auction market. The efficiency of a CDA market is established as being close to 100%. The author's model is validated in experimental settings, which correspond closely with those utilized in previous works in this area (Figs. 4.2, 4.3, 4.5). Further experimentation shows that it stays above 90% even when demand and supply are varied greatly (up to 50% change; see Figs. 4.6-4.8). Although previous research in agent based automated CDA markets has listed its many advantages with regards to its efficiency and property of quick convergence [Priest and Tol 1998; Priest and Merida-Campos 2002], the results presented here provide new data which establishes a quantitative and qualitative measure of this robustness (see Figs**.** 4.6-4.19) providing useful insight into the working of the CDA market in a dynamically changing marketplace where buyer seller populations keep changing.

The author believes that this work is an important step forward in our attempt to establish a tradeoff between various continuously clearing market mechanisms. The work presented is truly novel in such respects and will help fill a void[31] in the agent based automated markets community where the Double Auction has by far been the mechanism of choice for designers/researchers implementing/studying commodity markets.

---

31 In a personal communication with the author (Nov 2001); Dan Gode who first studied the structural properties of a CDA market by employing ZI traders [37], and Vernon Smith, who pioneered the field of experimental economics, allude to the absence of such a comparative study and its perceived usefulness.

This work is also novel in the sense that the author has modeled the CDA environment more realistically and in a more simplified manner. Previous researchers [Cliff 1997; Priest and Tol 1998] have modeled the Double auction as a synchronized bidding environment which actually translates to a Clearing House architecture (see Chapter 2 & 3) in the real world. As such, alluding to these models as a CDA environment is not entirely correct. This work is also different from previous research in that we allocate only a single item to each participant instead of the multiple goods associated with buyers/sellers in previous CDA implementations. This removes dubious auxiliary assumptions like multiple unit traders constructing strategies separately for each unit [Cason and Friedman 1995] and simplifies the experimental setup, providing the basis for a stark comparative assessment.

The experimental setup for a Quote Driven market in comparative settings is outlined. Again, the study of such a market mechanism for commodity markets has not been done previously. Work done with Quote Driven markets for securities trading [Chan and Shelton 2001] emphasizes the adaptive/learning process facing the market maker. The results presented for a CDA indicate that if demand/supply in the market is in excessive imbalance (ratio of buyers to sellers or vice versa is less than 1:2), the transaction prices drift away from those predicted by competitive equilibrium. Even when it is close to equilibrium, there is increased volatility around the equilibrium price.

5.2.2 Simulation Setup: Limitations

Some of the most restrictive assumptions in our experiments (these can be attributed to the experimentation that has followed on in the steps on previous researchers) include:

1) The experiments involve a stationary repetitive environment across different sessions of a market in which traders receive the same endowments each period. This does not hold true in realistic markets where demand and supply changes continuously across time. By simulating a CDA market with different activation rates for agents in each session, we have partly addressed this problem, however, the fact remains that the underlying demand/supply schedule remains the same through the period of the experiment. It would be interesting to explore the Double auction and Quote Driven markets in a random values environment (in which traders private values are drawn

independently each period from a known uniform distribution). This will help discard any aspects that might be attributable to selecting a particular demand/supply schedule[32].

2) As our aim is a comparative assessment of the two market mechanisms, we disregard any game theoretic behaviour on the part of the agents. Albeit, all our traders employ the same simplistic learning strategy. This is again not true in realistic markets. However, game theoretic analysis of markets with agents following different learning strategies is not simple. Most work to date has only attempted this in highly stylised settings. The recursive modelling method [Vidal and Durfee 1996] has been proposed as an approach for agents to reason about other agents. However in most practical cases the agent can build only finite nesting models due to the limitations of acquiring knowledge. Moreover, it has been shown that in an environment, where each agent is using a recursive modelling strategy to model its opponents, equilibrium is never reachable.

Both the above limitations can be addressed to some extent by utilising the real time trading platform that has been constructed.

5.2.3 Other Considerations

The Quote Driven market structure has been discussed in the previous chapters and an experimental scenario of bandwidth trading in a telecommunications network has been presented. It is clear from the discussion into the application of market mechanisms to perishable goods like electricity and bandwidth that it requires new thinking and concepts before it is close to realisation. Some of the relevant issues are:

1) *What commodities should be traded? (or equivalently the design of the product)*: Earlier, two types of trading scenarios have been alluded to i.e. the trading of simple commodities, which is conducted in a *spot market* and where agents have strictly private values for the goods. The other types of commodities are the standardized financial contracts as forwards, futures and even derivatives. In the Quote Driven market, the product is assumed to have

32 In Economic theory, some authors have considered particular demand supply schedules, which do not lead to efficient outcomes in a CDA.

either a common value or co-related value for the participants. For instance in the case of bandwidth trading, although the users might have different requirements which means that they assign different (or private) valuations to their requests, the market maker will need to be aware of these valuations generally because that impacts on the possibility of resale for him.

2) *What is the need for trading*? (*or equivalently what are the different types of traders in such a market*?): In a simple spot market for commodities, the sole interest of the participants is to derive material benefit from the act of exchange (i.e. money for a commodity). It is assumed that the commodity is consumed after purchase. In a Quote Driven market however, there can be several types of participants. One class of them is similar to those in simple commodity markets i.e. those who trade for exogenous reasons. Since the commodity is long lived (i.e. there is a possibility of resale), there might be other participants who act as speculators or arbitrageurs (people who hope to make risk-less profits).

5.3 Further Work

The work presented here has established the framework for a comparative assessment of the Double Auction mechanism against a novel automated market mechanism in the Quote Driven market. The author believes that this could provide new results, which will help establish the tradeoffs with a Double Auction market in an empirical setting.

The thesis has also outlined possible applications of the Quote Driven market in specific domains. For bandwidth trading, a Quote Driven market for fixed line telecommunication networks has been presented. This will need to be developed further within a test bed scenario. Empirical evaluation could provide an insight into generic issues like price volatility, liquidity, security etc. besides domain specific metrics like call set-up delay, possibility of disconnection etc. Extending the application of the business scenario involving the Quote Driven market to a wireless domain offers a much more realistic scenario. This is because in wireless networks[33] capacity constraints, higher bit rate services, and differing business objectives of

33 The emergence of **M**obile **V**irtual **N**etwork **O**perators in wireless networks is indicative of the emergence of a multi SP model.

SPs/NPs and users necessitate constant negotiation for resources[34]. Although the studies on bandwidth markets presented previously have mostly targeted fixed networks, bandwidth (capacity) trading in a wireless network is more likely because the radio resources (including capacity) are scarce and are insufficient for a particular SP to reserve independent capacity in advance; something which can assumed true in wired networks.

The figure below indicates a situation where a service provider's resources (available bandwidth) are overstretched for a short 'peak' period. In normal circumstances this would mean calls arriving after C4 will be blocked. However, if we envisage a market where service providers can negotiate to lease bandwidth to accommodate such peaks, it will lead to better service provision and more efficient utilisation of network resources. It is the author's belief that implementing novel allocation/negotiation mechanisms that allow the development of new business models to support such a scenario will go a long way to bringing this to fruition.

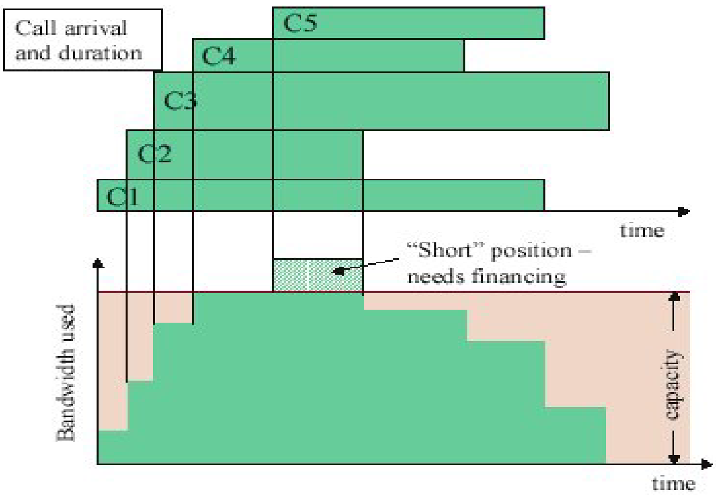

(**Fig**. 5.1 Market making in telecommunication capacity[35])

5.4 Other possible areas of exploration

1) Analysis of improved learning mechanisms to be incorporated into the behaviour of the trading agents.

---

34 http://www.elec.qmul.ac.uk/research/projects/shuffle.html

35 Courtesy EU IST project SHUFFLE

2) Use of the Quote Driven market mechanism in distributed resource allocation/optimisation schemes i.e. a closed market as a control mechanism in suitable areas within telecommunications.
3) Comparative assessment of the Quote Driven market with other market protocols for example sealed bid auctions, Vickrey auctions, reverse auctions etc.
4) Research into multi-attribute negotiation within a CDA and Quote Driven market environment.

5.5 In Conclusion

The work outlined here presents a novel assessment of the CDA mechanism and a new Quote Driven approach for automated markets that can potentially be used to develop more liquid and accessible market places. The research argues that providing incentives to third party traders (market makers in a Quote Driven market) can lead to a better overall service for clients wishing to trade goods in real time. This work has established a specification of market roles, protocols and infrastructure as well as suitable evaluation metrics for a comparative assessment of automated electronic markets [Bourne and Zaidi 2001] which could be extended with further research.

Emerging commodities such as bandwidth, energy, computation etc. can be provisioned in real time and will simply expire if not used. This calls for a new kind of market for real time trading in such 'resource goods' in which software agents can value resources and participate in real-time negotiation to share them. The author believes that the Quote Driven model offers a promising mechanism for supporting such environments. Although the thesis offers an exposition of bandwidth trading, the model is applicable to other commodities including electric power and so forth.

The work has led to the development of a discrete time simulator, which is being utilised for a detailed assessment of the Double auction market. The results to date have validated the model for a CDA market and have also provided concrete quantitative evidence of the robustness of a CDA market under changing demand supply conditions.

References[36]:

Bailey, J. and Y. Bakos (1997). "An Exploratory Study of the Emerging Role of Electronic Intermediaries". International Journal of Electronic Commerce **1**(3): 7-20.

Bigham, J., L. G. Cuthbert, et al. (1999). "A Distributed and Decentralised Architecture for Control of Connection Admission to a Telecommunications Network". Proceedings of the IMPACT Conference: (5-15). London.

Bourne, R. (2000). "A Quote-Driven Market for Service Providers in Telecommunications Networks". Proceedings of the 8th International Conference on Telecommunication Systems Modeling and Analysis: (205-209). Nashville, TN, USA.

Bourne, R. and R. Zaidi (2001). "A Quote-Driven Automated Market". Proceedings of the Symposium on Information Agents for E-Commerce at the AISB Convention. York.
http://www.elec.qmul.ac.uk/people/rach/documents/aisb01_000.pdf

Bower, J. and D. Bunn (1999). "A Model Based Comparison of Pool and Bilateral Market Mechanisms for Electricity Trading". Technical Report, London Business School, Energy Markets Group.
http://www.econ.iastate.edu/tesfatsi/retasim.doc

Cason, T. N. and D. Friedman (1995). "Price formation in double auction markets". Journal of Economic Dynamics and Control **20**: 1307-1337.

Chan, N. and C. Shelton (2001). "An Adaptive Electronic Market-maker". Proceedings of the 7th International Conference of the Society for Computational Economics. Yale University.
http://www.cs.ucr.edu/~cshelton/papers/AIM-2001-005.pdf

Chao, H.-p. and R. Wilson (1999). "Design of Wholesale Electricity Markets". Electric Power Research Institute.
http://faculty-gsb.stanford.edu/wilson/archive/E542/classfiles/draft990106.pdf

Cheliotis, G. (1999). "A Market-Based Model of Bandwidth and a New Approach to End-to-End Path Computation with Quality Guarantees". Workshop on Internet Service Quality Economics, MIT. Cambridge, MA, USA. http://www.zurich.ibm.com/pdf/bandwidth/bandwidth_end-to-end_path.pdf

Cheliotis, G. (2000). "Bandwidth Trading in the Real World: Findings and Implications for Commodities Brokerage". IBM Research. http://www.cis.upenn.edu/markets/ReadingGroup/meetings/meetings.htm

[36] In some instances the referenced Conference/Workshop Proceedings items have been sourced online. In such cases, no overall numbering scheme is apparent vis-à-vis the published Proceedings.

Cheng, J. Q. and M. P. Wellman (1998). "The WALRAS algorithm: A convergent distributed implementation of general equilibrium outcomes". Computational Economics **12**: 1-24.

Clearwater, S. H., Ed. (1996). "Market Based Control: A paradigm for Distributed Resource Allocation", World Scientific Press.

Cliff, D. (1997). "Minimal-Intelligence Agents for Bargaining Behaviors in Market-Based Environments". Technical Report, HP Labs. http://www.hpl.hp.com/agents/papers.htm

Cliff, D. and J. Bruten (1997). "Zero is Not Enough: On the Lower Limit of Agent Intelligence for Continuous Double Auction Markets". Technical Report, HP Labs. http://www.hpl.hp.com/agents/papers.htm

Cohen, K., S. Maier, et al. (1986). "The Microstructure of Securities Markets", Prentice Hall.

Courcoubetis, C., A. Dimakis, et al. (2001). "Providing Bandwidth Guarantees over a Best-effort Network: Call-admission and Pricing". Proceedings of INFCOM 2001: (459-467). Anchorage, Alaska.

Courcoubetis, C. and R. Weber (2003). "Pricing Communication Networks: Economics, Technology and Modelling", Wiley Europe.

Das, R., J. Hanson, et al. (2001). "Agent Human Interactions in the Continuous Double Auction". Proceedings of IJCAI: (1169-1187). Seattle, USA.

Davis, D. and C. Holt (1993). "Experimental Economics", Princeton University Press.

Economides, N. (2004). "Economics of Networks". International Journal of Industrial Organization **14**(6): 673-700.

Faratin, P. (2000). "Automated Service Negotiation Between Autonomous Computational Agents". PhD Thesis, Department of Electronic Engineering,Queen Mary, University of London.

Flood, M., R. Huisman, et al. (1999). "Quote Disclosure and Price Discovery in Multiple-Dealer Financial Markets". The Review of Financial Studies **12**(1): 37-59.

Foss, J., K. Garcha, et al. (2000). "Brokerage in an Information Economy". Proceedings of The 10th Annual Internet Society Conference, Yokohama, Japan. http://www.isoc.org/inet2000/cdproceedings/7a/7a_1.htm

Friedman, D. and J. Rust, Eds. (1992). "The Double Auction Market: Institutions Theories and Evidence". New York, Addison Wesley.

Fudenberg, D. and J. Tirole (1991). "Game Theory", MIT Press.

Garman, M. (1976). "Market Microstructure". Journal of Financial Economics **3**: 257-175.

Gibbens, R. and F. Kelly (1999). "Resource Pricing and Evolution of Congestion Control". Automatica **35**: 1969-1985.

Gibney, M. A., N. R. Jennings, et al. (1999). "Market-Based Call routing in Telecommunication Networks". Proceedings of The Third International Workshop on Multi-Agent Systems and Telecommunications: (46-61). Stockholm, Sweden.

Gjerstad, S. and J. Dickhaut (1998). "Price Formation in Double Auctions". Games and Economic Behaviour **22**: 1-29.

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

Grigonis, R. (2001). "Bandwidth for the Buck". Computer Telephony, http://www.cconvergence.com/

He, M., H.-f. Leung, et al. (2002). "A Fuzzy Logic based bidding strategy for autonomous agents in Continuous Double Auctions". IEEE Transactions on Knowledge and Data Engineering **5**(6): 1345-1363.

Huberman, B., Ed. (1988). "The Ecology of Computation", Elsevier/North Holland.

IBM-Research: "Information Economies Project at IBM". http://www.research.ibm.com/infoecon/

Jennings, N. R., P. Faratin, et al. (1996). "Using Intelligent Agents to Manage Business Processes". First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology: (345-360).

Katz, R. H., I. Stoica, et al. (2002). "The SAHARA Model for Service Composition across Multiple Providers". International Conference on Pervasive Computing,(1-14).

Kelly, F. (2001). "Mathematical Modelling of the Internet". Mathematics Unlimited-2001 and Beyond: (685-702). Eds. B. Engquist and W. Schmid. Berlin, Springer-Verlag.

Kephart, J., J. Hanson, et al. (2000). "Dynamic Pricing by Software Agents". Computer Networks **32**: 731-752.

Kersten, G. E. and G. Lo (2001). "Negotiation Support Systems and Software Agents in E-business Negotiations". http://interneg.concordia.ca/interneg/research/papers/2001/05.pdf

Klemperer, P. (2002). "How (not) to run auctions: The European 3G telecom auctions". European Economic Review **4-5**(46): 829-845.

Kollock, P. (2000). "The World Wakes Up To Online Exchanges". @Markets Magazine Premier Issue: 12-17.
Kreps, D., M. (1990). "A Course in Microeconomic Theory", Harvester Wheatsheaf.

Lamper, D. and N. F. Johnson (2002). "The Science of Complexity: Finding out when two's company and three's a crowd." Dr. Dobb's journal Intelligent Systems 341: 16-22.

Lehr, W. and L. McKnight (1998). "Next Generation Internet Bandwidth Markets". Communications and Strategies, Special Issue on Internet Economics **32**: (91-106).

Lukose, R. and B. Huberman (1997). "A Methodology for Managing Risk in Electronic Transactions over the Internet". Technical Report HP Labs. http://www.hpl.hp.com/research/idl/projects/ecommerce/banking.ps

Madhavan, A. (1992). "Trading Mechanisms in Securities Markets". Journal of Finance **47**: 607-642.

Madhavan, A. (2000). "Market Microstructure: A Survey". Journal of Financial Markets **3**(3): 205-258.

Mas-Colell, A. (1995). "Microeconomic Theory". New York, Oxford University Press.

McAfee, P. and J. McMillan (1996). " Analyzing the Airwaves Auction". Journal of Economic Perspectives **10**(1): 159-175.

McMillan, J. (1994). "Selling Spectrum Rights". Journal of Economic Perspectives **8**(1): 145-162.

Mindel, J. and M. Sirbu (2001). "Taxonomy of Traded Bandwidth". Technical Report, Engineering and Public Policy Department, Carnegie Mellon University.

Mitchell, P. and J. Moore (2001). "Growing the MVNO Model - A goal shared by all parties". EMC Research.
http://adventis.com/mvno/pdf/EMC_brief.pdf

OFTEL (2002). "3G Infrastructure sharing."
http://www.oftel.gov.uk/publications/mobile/index.htm

O'Hara, M. (1995). "Market Microstructure Theory", Blackwell Publishers.

Park, S., E. Durfee, et al. (1999). "An Adaptive agent bidding strategy based on stochastic modelling". Proceedings of Third International Conference on Autonomous Agents: (147-153).

Priest, C. and C. M. Merida-Campos (2002). "Automated Negotiation in Many-to-Many markets for Imperfectly substitutable goods". Technical report, HP Labs. http://www.hpl.hp.com/techreports/2002/HPL-2002-197.pdf

Priest, C. and M. v. Tol (1998). "Adaptive Agents in a Persistent Shout Double Auction". Proceedings of The International Conference on Internet Computing and Economics: (11-18), ACM Press.

Rasmusen, E. (2001). "Games and Information", Blackwell Publishers.

Rasmusson, L. (2001). "Pricing Virtual Paths with Quality-of-Service Guarantees as Bundle Derivatives". Technical Report, SICS, Stockholm, Sweden. http://xxx.lanl.gov/abs/cs.NI/0106028

Roth, A. (2002). "The Economist as an Engineer: GameTheory, Experimental Economics and Computation as tools of Design Economics". Econometrica **70**: 1341-1378.

Rust, J., J. Miller, et al. (1992). "Behaviour of Trading Automata in a computerized Double Auction Market". The Double Auction Market: Institutions Theories and Evidence. D. Friedman and J. Rust. New York, Addison Wesley.

Sandholm, T. (1999). "Distributed Rational Decision Making". In "Multiagent Systems". Ed. G. Weiss, MIT Press.

Sandholm, T. (2000). "eMediator: A Next Generation Electronic Commerce Server". Proceedings of The International Conference on Autonomous Agents (AGENTS): (341-348). Barcelona, Spain.

SEBI (2000). "Report of the committee on market making". Securities and Exchange Board of India. www.sebi.com

Semret, N. and A. Lazar (1999). "Spot and Derivative Markets in Admission Control". Proceedings of the ITC 16: (925-941). Edinburgh, UK, Elsevier.

Smith, V. (1962). "An Experimental study of Competitive Market Behaviour". Journal of Political Economy **70**: 111-137.

Smith, V. (1982). "Microeconomic Systems as an Experimental Science". The American Economic Review **72**(5): 923-955.

Stoll, H. (1978). "The Supply of Dealer Services in Securities Markets". Journal of Finance **33**: 1133-1151.

Tesauro, G. and R. Das (2001). "High-Performance Bidding Agents for the Continuous Double Auction". Proceedings of The IJCAI Workshop on Economic Agents, Models, and Mechanisms. Seattle, USA. http://www.research.ibm.com/infoecon/paps/dblauc.pdf

Theissen, E. (2000). "Market Structure, informational efficiency and liquidity: An experimental comparison of auction and dealer markets". Journal of Financial Markets(3): 333-363.

Varian, H. (2004). "Network Economics". International Journal of Industrial Organization **14**(6): 673-700.

Vidal, J. and E. Durfee (1996). "The impact of nested agent models in an information economy". Proceedings of The Second International Conference on Multi agent systems: 377-384.

Vulkan, N. and C. Priest (1999). "Automated trading in agent based markets for communication bandwidth". Technical Report, HP Labs.
http://www.hpl.hp.com/techreports/2000/HPL-2000-24.pdf

Wellman, M. (1996). "Market-oriented programming: Some early lessons". Market-Based Control: A Paradigm for Distributed Resource Allocation. S. Clearwater, World Scientific Publishing Co.

Wellman, M. P., W. E. Walsh, et al. (2001). "Auction protocols for decentralized scheduling". GAMES AND ECONOMIC BEHAVIOR. **35**: 271-303

Wilson, R. (1999). "Market Architecture". Technical Report, Stanford University.
http://www.market-design.com/files/wilson-market-architecture.pdf

Wolak, F. (1997). "Market Design and Price Behavior in Restructured Electricity Markets:An International Comparison" Technical Report, Stanford University. ftp://zia.stanford.edu/pub/papers/taiwan.pdf

Wurman, P. (2001). "Dynamic Pricing in the Virtual Marketplace". IEEE Internet Computing. **5**(2): (36-42) Mar-Apr

Ygge, F. (1998). "Market-Oriented Programming and its Application to Power Load Management". PhD Thesis, Department of Computer Science, Lund University.

Ygge, F. (1999). "Energy Resellers – An Endangered Species?" Technical Report, EnerSearch AB, Sweden.
http://www.enersearch.se/~ygge/

Ygge, F. and H. Akkermans (1999). "Decentralized Markets versus Central Control: A Comparative Study". Journal of Artificial Intelligence Research (JAIR) **11**: 301-333.

Zaidi, R. (2001). "Quote-Based Automated Markets (Stage 1 Report)". Queen Mary, University of London.

Zaidi, R., R. Bourne, et al. (2002). "Automated Markets: Intermediation and Market Making". Proceedings of PGNET: (86-91), Liverpool.