EVOLUTIONARY MARKET-BASED RESOURCE ALLOCATION IN DECENTRALISED COMPUTATIONAL SYSTEMS

by

PETER RICHARD LEWIS

A thesis submitted to
The University of Birmingham
for the degree of
DOCTOR OF PHILOSOPHY

School of Computer Science
College of Engineering and Physical Sciences
The University of Birmingham
September 2010

UNIVERSITY OF BIRMINGHAM

University of Birmingham Research Archive

e-theses repository

This unpublished thesis/dissertation is copyright of the author and/or third parties. The intellectual property rights of the author or third parties in respect of this work are as defined by The Copyright Designs and Patents Act 1988 or as modified by any successor legislation.

Any use made of information contained in this thesis/dissertation must be in accordance with that legislation and must be properly acknowledged. Further distribution or reproduction in any format is prohibited without the permission of the copyright holder.

Abstract

This thesis presents a novel market-based method, inspired by retail markets for resource allocation in fully decentralised computational systems where agents are self-interested. The posted offer mechanism used requires no central or regional coordinator or complex negotiation strategies. The stability of outcome allocations, those at equilibrium, is analysed and compared for three buyer behaviour models. The approach is scalable, robust and may be tuned to achieve a range of desired outcome resource allocations. These include a balanced load, allocations reflective of providers' differing capabilities and those appropriate to heterogeneous buyer preferences over multiple attributes.

The behaviour of the approach is studied both game theoretically and in simulation, where novel *evolutionary market agents* act on behalf of resource providing nodes to adaptively price their resources over time in response to market conditions. Sellers competitively co-evolve their offers online without any need for global market information. This is shown to lead the system to the game theoretically predicted outcome resource allocation when buyers' decision functions degrade gracefully. Additionally, allocations remain stable in the presence of small changes in price and other more disruptive agents. The posted offer model therefore appears to be a useful mechanism for resource allocation in both homogeneous and heterogeneous decentralised computational systems where nodes are self-interested. Furthermore, evolutionary computation is shown to be a potential approach to realising self-interested adaptive pricing behaviour under the assumption of private information present in the posted offer model.

For all those from whom I have learnt,

especially the previous generation.

Acknowledgements

I would like to give my sincere thanks to Professor Xin Yao for being an excellent guide through this fascinating journey, continually providing wisdom and genuinely constructive criticism. He has let me indulge myself and explore new crazy ideas on my own, judging the balance perfectly between pointing out the mountain tops and letting me discover the route for myself. For this, I am very grateful.

To Dr Paul Marrow I would also like to say thanks. Paul has provided a valuable external perspective throughout, keeping my work relevant and my assumptions grounded. His and BT's support throughout this project have provided numerous exciting and challenging opportunities, and I am delighted to have been able to take advantage of them.

Also I am grateful to the many others who have helped me along the way over the last few years. Particularly, my thanks go to Dr Jon Rowe, Dr John Bullinaria and Dr Alastair Channon who took the time to provide guidance as members of my Thesis Group, for their regular critical eye and words of insight and perspective. Thanks also to Chris Staite, Damien Duff and Dr Martín Escardó, for lending me time with their greater knowledge of mathematics. Additionally, thanks to those involved in the EPSRC Market-Based Control of Complex Computational Systems project for their regular feedback. Greatest amongst these was Dr Peter McBurney.

Warm thanks go to my office mates over the last four years: Seyyed Shah, Thomas Miconi, Nur Hana Samsudin and Mohammed Alodib. They made the daily business of being a research student a real pleasure and I'm going to miss sharing *the best office on campus* with them and the yukka.

I am greatly indebted to Edd Robinson, Arjun Chandra and Vivek Nallur, with whom I have worked most closely. Their contributions and participation in endless afternoons with the whiteboard really made the difference at key points over the last few years. Working with good friends is a privilege not many get to enjoy.

Indeed, thanks also go to all those teachers who have helped me get to where I am today: to Mike Boucher for teaching me to ask why, and that there's always another question to be asked; to Dave Freeman for encouraging me to go beyond simply what's needed; to David Hasdell for helping me realise that I love to write; and to Rick Thomas for getting me excited about Computer Science.

Gratitude and happiness beyond measure go to my beautiful future wife Marla, for being so amazingly understanding of the seemingly random ups and downs of the last few years. She has demonstrated time and again that she really gets why I wanted to do this whole thing anyway, and her words and actions have helped me succeed in so many ways.

Immense thanks go also to my family: to my brother Dave, for sharing my vision that you can go your own way; to my Nan for showing me that anything is possible once you've decided to do it; to my Mom for being a tremendous inspiration and showing me what success looks like; and to my Dad for never failing to present a semi-serious yet incredibly wise and motivating perspective on life, just at the right moment. Indeed more thanks than I can hope to express go to both my parents. They taught me the value of striving to fulfil my potential, which led me to never be satisfied with having learnt enough. I'm not done yet.

Publications arising from this Thesis

- [1] P. R. Lewis, P. Marrow, and X. Yao, "Evolutionary market agents for resource allocation in decentralised systems," in *Parallel Problem Solving From Nature - PPSN X*, vol. 5199 of *Lecture Notes in Computer Science*, pp. 1071–1080, Springer, 2008
- [2] P. R. Lewis, P. Marrow, and X. Yao, "Evolutionary market agents and heterogeneous service providers: Achieving desired resource allocations," in *Proceedings of the IEEE Congress on Evolutionary Computation*, pp. 904–910, 2009
- [3] P. R. Lewis, P. Marrow, and X. Yao, "Resource allocation in decentralised computational systems: An evolutionary market based approach," *Autonomous Agents and Multi-Agent Systems*, vol. 21, no. 1, pp. 143–171, 2010

Patent arising from this Thesis

- [4] P. R. Lewis, P. Marrow, and X. Yao (named inventors), "Load balancing in a data network." European Patent Application No. 08252953.8, filed September 2008. UK Patent Application No. 0905439.6, filed March 2009. US / Canadian Patent Application No. PCT/GB2009/002135, filed September 2009. Korean Patent Application No. A31728 KRw, filed September 2009. Holder: BT PLC

Contents

List of Figures

CHAPTER 1

INTRODUCTION AND MOTIVATION

One scenario of the future of computation populates the Internet with vast numbers of software agents providing, trading, and using a rich variety of information goods and services in an open, free-market economy.

Jeffrey Kephart, James Hanson and Jakka Sairamesh

The environment within which information and computation systems operate is changing. High bandwidth always-on communication channels are everywhere and telecommunications infrastructure is continuing to be deregulated. This leads to opportunities for the creation of vastly scalable computational platforms, built from generic and substitutable components. As a result, organisations and individuals are increasingly finding themselves operating in scenarios where they do not have full control or even full knowledge of key resources such as network infrastructure, data provision and computational capacity. These are resources which nonetheless remain vital to their continued successful operation and development.

In this context, paradigms are being created for the development and deployment of massively distributed computational and business systems, which enable resources to span many locations, organisations and platforms, connected through the Internet. Grid [5], cloud [6] and service oriented [7] computing are examples of such paradigms. Due to their flexibility and scalability, approaches which construct systems from generic distributed components are becoming highly prominent in both distributed computing and e-commerce [8].

In order for any organisation to work successfully within such a computational ecosystem, there is a need to find novel ways to manage, control and understand these highly complex distributed systems. Research has now begun to look towards socially and economically inspired agent-based systems to manage the allocation of distributed components in a necessarily dynamic environment [6, 7]. This thesis aims to contribute to this approach.

1.1 Scenario

Future distributed, decentralised computational paradigms envisage vast numbers of software agents, representing resource providers and users as they arrive, organise and dissipate, and computational capabilities are formed and reformed as needed, without reference to a central authority or coordinator. It is predicted [9] that in supporting this, the Internet will evolve into a space in which these agents will engage in a market economy, trading for a wide range of generic computational resources. Indeed, as the Internet matures, it is predicted that the majority of traffic it carries will be generated by autonomous agents, acting on behalf of their owners [10]. Such distributed systems are considered in this thesis in the broadest sense, as systems where there exists an amount of work or resource to be distributed about a network of nodes. However, no single agent is expected to have either full control or knowledge of the system as a whole. As these systems continue to be developed, there is a need to find novel ways to understand and autonomically manage and control them [11]. As part of this, there remains the problem of how such resources can or should be allocated amongst the nodes [12].

From an engineering perspective, any resource allocation problem will have a set of objectives: a desired allocation or outcome. A common example of this is a balanced load, where the task of providing a resource is shared evenly between a group of nodes. More complex objectives may involve stable, uneven allocations, where account is taken of other factors. Such factors might include users' preferences over quality of service issues, underlying costs to the resource provider, or differences in the capability of nodes to provide an equivalent resource. The ability to allocate resources in a desired configuration, in a scalable and robust manner, is essential.

In many systems, nodes are either assumed or designed to be *self-interested*, each wishing for example to maximise its allocation or perform its assigned task as fast as possible without regard for other tasks. In these cases, either the total demand for a resource may exceed its total supply or else the unrestricted use of resources regardless of task prioritisation can negatively affect overall performance. The allocation of resources to individuals in the presence of such scarcity is the classic problem studied by economics. In the human world, building on the simple idea of bilateral exchange, it is often approached using the structure and rules of a market, and its tool, price.

Markets can also be used in settings where it is impractical or unwanted to associate monetary payment with the resource allocation process. Indeed, market-based control is a broad approach to resource allocation in a wide range of real world applications [13]. In these scenarios, artificially created markets and pricing can enable stable, efficient, self-organising resource allocation to be achieved. Where resources can be treated as interchangeable commodities able to be provided by anyone, market economies can be used to allocate such resources when individuals' decisions are self-interested [14].

However, the market mechanisms employed in these systems to date, though operating in a distributed manner, often rely upon some form of central or regional coordination or control, such as an auctioneer, specialist, or set of super-nodes [15]. It is likely that this lack of full decentralisation leads to unfulfilled potential in terms of both scalability and robustness to failure. Alternatives, such as those approaches which rely on agents' ability to negotiate bilaterally are promising options, however often rely on the complex cognitive ability of agents or highly developed strategies. This will add an additional computational overhead at each node, which may be avoidable. This thesis considers a further less studied alternative, that of the retail-inspired posted offer market mechanism. The posted offer mechanism was first described by economists studying large retail stores, in which prices are fixed by the sellers and publicly announced. Buyers subsequently choose how much to buy, and from whom, based on these posted prices. The mechanism is used widely in human markets, especially where large numbers of buyers value resources similarly. It also has the advantages of being simple to understand and implement, as well as requiring little computational overhead on the part of participants.

1.2 Overarching Research Questions

This thesis is concerned with two overarching research questions:

- How do posted offer markets allocate resources in decentralised computational systems consisting of self-interested individuals? More specifically, what can we say about what outcomes to expect from employing them and what gaps, if any exist between ideal cases and reality?
- How can we design software agents which perform well in such systems, in order to achieve their objectives and in doing so give rise to the predicted global outcomes?

These questions are studied in an abstract model of a decentralised computational system consisting of resource providing and resource using nodes. Analysis is performed with a view to obtaining desired global outcomes, brought about by the interactions between the self-interested nodes. Dependent on the behaviour and properties of the nodes, interesting and useful outcomes can be predicted. These include the load being evenly balanced across the network, resources allocated according to nodes' capabilities and resources being provided with attributes appropriate to the preferences and demands of the users. Throughout, the impact of different and heterogeneous node behaviours and capabilities are considered, and experimental results additionally provide evidence of robustness and scalability.

Using a range of example scenarios, obtainable outcome resource allocations are described game theoretically, with reference to classic economic models such as Bertrand competition. Furthermore *evolutionary market agents* are developed, which enable the demonstration of these outcome allocations using evolutionary simulations of agents' interactions. Unlike the game theoretic analyses, these simulations replicate more accurately the assumptions of incomplete and private information present in real world embodiments of the model. Where differences exist between the outcomes in the theoretical and simulation results, analysis is provided to explain why this is the case, and what if anything may be done to mitigate against it.

This exciting topic is by its very nature interdisciplinary. Whilst the application area is firmly rooted in computing, the processes and techniques involved draw heavily from economics and evolutionary biology. An introduction to relevant concepts is also therefore provided.

1.3 Contributions of the Thesis

The major contributions of this thesis are as follows:

- The description of a novel method of allocating resources in fully decentralised computational systems using self-interested agents. The approach does not rely on any central or regional coordinating node, nor complex strategic reasoning on the part of participating agents.

- A method for and examples of game theoretic analysis of the mechanism, describing how equilibria and hence resulting outcome resource allocations may be predicted, given a particular population of agents.

- A description of how to use the approach in order to obtain an evenly balanced load across a decentralised network of resource providing nodes, in the absence of central coordination or cooperation.

- An analytical method for determining parameters for the system in order to bring about additional desired global outcome resource allocations.

- A novel agent strategy, which makes use of evolutionary computation to adaptively price resources over time and is capable of achieving the predicted outcome allocations using only private information.

- An extension to the approach, grounded in decision theory, which permits resources to be described over multiple quality attributes. Analysis demonstrating that the extended approach leads to resources being provided with quality of service attributes determined appropriately and proportionally to a heterogeneous population of users.

- A comparison of the performance of the mechanism when user agents are endowed with a range of different decision making behaviours.

As was described in publications arising from this thesis [1, 2, 3], the method presented here joins only two other families of market-based approaches for achieving resource allocation under the assumption of self-interest in fully decentralised systems with no central or regional coordination. The alternatives, either using bilateral negotiation or provider nodes hosting their own auctions, require a larger degree of complex strategic reasoning, particularly on the part of resource users' agents in order to perform well. The approach presented in this thesis does not require users to use such complex agents; considerably simpler strategies are shown to be sufficient in many cases.

Additionally, the line of questioning taken in this thesis departs from those employed in developing traditional market-based control methods, in that it focuses on the achievement of particular stable outcome resource allocations as the desired objective. Three particular types of objective are studied:

- A balanced load, where each node provides an equivalent quantity of the resource to its users;
- Additional stable uneven allocations, where the quantity provided by each node may be tuned according to parameters; and
- Resources with heterogeneous quality attributes provided such that the quality attributes are determined automatically in order to suit the demands of a heterogeneous population of users.

Throughout the thesis, a methodological approach inspired by Kephart et al. [9] is taken. For each type of objective, firstly the behaviour of *myoptimal* agents is investigated, that is those able to accurately optimise their behaviour based on the current state of the environment. Subsequently the key assumption required for *myoptimal* behaviour, of full knowledge of the current state of the environment is relaxed, and agents instead take an evolutionary approach, making use of only private information. The resulting behaviour of each is compared and commonalities and differences are highlighted. As will be shown, the first approach is typically only feasible in very small systems, where the environment is sufficiently simple to allow an agent's behaviour to be optimised in real time with respect to its expected payoff. Since the evolutionary approach makes use of only private information, such global optimisation of behaviour is not required. Therefore much larger and more complex scenarios are able to be considered in this case.

1.4 Overview of the Thesis

The remainder of this thesis is structured as follows. Chapter 2 introduces the problem of resource allocation in decentralised computational systems, and discusses why parallels can be drawn between such systems and social and economic ones. A number of existing classical and market-based approaches are reviewed, highlighting why the development of a new approach is valuable for use in systems with self-interested participants and no central control or coordination. In chapter 3 a resource allocation model, the problem to be studied is formulated, based on the motivating scenario. The fully decentralised market-based approach is described here and user behaviours to be investigated are also introduced. Initial game theoretic analysis of the expected behaviour of the system is presented for each user behaviour, illustrating how a balanced load may be achieved when provider agents are *myoptimal*. Chapter 4 relaxes the assumption of complete knowledge required by *myoptimal* agents and instead presents the *evolutionary market agent* algorithm, which requires only private information. The analysis is then extended to examine the impact of this on the system, and a load balanced outcome is demonstrated in simulation. Results on scalability and mixed buyer populations are also presented in this context. Chapter 5 extends the model to consider heterogeneous provider node capabilities by taking into account differing costs of resource provision and nodes' private valuations. Both the analytical method from chapter 3 and the evolutionary simulations from chapter 4 are employed in order to predict and design for a variety of additional outcome resource allocations. Chapters 6 and 7 consider a further and highly realistic extension to the model, which accounts for resources that are described over multiple quality attributes. A description of how agents' decision making processes can also be extended to account for this is presented, based on multi-attribute utility theory. This extension provides a qualitative difference in terms of the difficulty of the problem studied, as competition is no longer conducted on price alone. Game theoretic analysis and evolutionary simulations are again employed in order to predict and demonstrate stable outcome resource allocations, under a variety of behavioural conditions. Additionally, the simulation results in chapter 7 highlight a dilemma faced by sellers represented by *evolutionary market agents*, brought about by the additional complexity of the multi-attribute extension. The dilemma is characterised through further experimental results, and the impact of this on overall system performance is quantified. Finally, chapter 8 concludes the thesis by reviewing the contributions of the preceding chapters and discussing future prospects and directions for further research.

CHAPTER 2

THE ECONOMICS OF COMPUTATIONAL RESOURCE ALLOCATION

Some of the owner men were kind because they hated what they had to do, and some of them were angry because they hated to be cruel, and some of them were cold because they had long ago found that one could not be an owner unless one were cold. And all of them were caught in something larger than themselves. Some of them hated the mathematics that drove them, and some were afraid, and some worshipped the mathematics because it provided a refuge from thought and from feeling. If a bank or finance company owned the land, the owner man said: The Bank - or the Company - needs - wants - insists - must have - as though the Bank or the Company were a monster, with thought and feeling, which had ensnared them. These last would take no responsibility for the banks or the companies because they were men and slaves, while the banks were machines and masters all at the same time. Some of the owner men were a little proud to be slaves to such cold and powerful masters.

John Steinbeck

This chapter provides an introduction to the problem of computational resource allocation and a review of previous solutions. After identifying common ground between decentralised computational systems and economic systems, the application of economic principles to the allocation of resources in such systems is considered and existing approaches discussed. In particular, two broad families of market-based control approaches are identified, those which use auctions and those which rely on bilateral bargaining. A third promising but less well understood approach is described, based on retail markets, and reasons are highlighted why, from the perspective of simplicity this is an attractive option. The small amount of existing work applying this mechanism to resource allocation is discussed, along with some prior methodology, specifically the use of evolutionary computation and game theoretic analyses. Considerable scope is shown for research in applying these techniques to better understand how the approach may be deployed in solving computational resource allocation problems.

This thesis in general and this chapter in particular assume no prior expert knowledge of economics, and as such an accessible introduction to relevant concepts is provided. Additionally, a brief introduction is given to other key concepts relied upon in the thesis, including game theory, agents and evolutionary computation.

This chapter proceeds as follows. Section 2.1 introduces and motivates the family of computational resource allocation problems, with reference to relevant technologies including grid, cloud and service oriented computing. Previous centralised solutions are described, along with issues which arise from the use of centralisation and potential decentralised alternatives. A discussion is also presented in this section on the differences between cooperative, non-cooperative and self-interested node behaviour. Section 2.2 describes why the systems in which this thesis is interested may, under certain conditions, be thought of as having much in common with economies; parallels are drawn between resource allocation in computational systems and the human world. This section also introduces some important concepts from economics which are relied upon in subsequent chapters, including game theory. Section 2.3 looks in more detail at how economics has previously been applied to computational resource allocation and reviews the families of approaches which currently exist. A case is made for the further study of the applicability and impact of the posted offer market mechanism. Section 2.4 discusses a particular problem to emerge from the use of the posted offer model, that of adaptive pricing by sellers, and again existing approaches to this problem are discussed. Finally, section 2.5 concludes with a summary of the chapter, drawing conclusions and highlighting why the work presented in this thesis is needed.

2.1 Computational Resource Allocation

Since the earliest days of computers, people have sought to apply them to the solving of large and complex problems. Indeed, computers' ability to solve large problems have brought benefits to humanity in fields as wide ranging as chess playing [16] and protein folding [17], amongst many others. However, key to the continued ability to apply computers to these kinds of problems is finding ways to enable them to scale massively, while remaining accessible to those who might use them. For example, it could be argued that the requirement either to own a supercomputer such as Deep Blue or have the funds and specialist knowledge to build a distributed platform such as that used by the Protein@Home project reduces accessibility.

Grid computing is one technology which attempts to address this. By providing a standard way to access computing power *on tap*, a grid platform allows users to run very large generic programs, distributed over many computational nodes [5]. Related technologies such as cloud computing [6] enable a similar standard means of access to potentially unbounded scalable computing, while service oriented architectures [7] provide a framework for distributed computational resources to be componentised and packaged up, such that distributed applications may be constructed from loosely coupled components.

Given such a range of approaches to scaling up computational capabilities, it is not surprising that computational resource allocation in such systems is not a single well defined problem. Instead, it is perhaps best described as a family of problems, each specific to the particular embodiment, but with much in common. At its heart however, the problem of computational resource allocation can be stated as follows: how should computational resources be made available to users, such as to achieve the objectives of the resource providing nodes, the users and the system overall? It is important to note here that the term *user* does not apply solely to an end user of a computer system or their processes, but also to any component which requires the use of a resource from a providing node.

In order to answer this question for a particular system, it is of course necessary to possess some further understanding of what is required. Does the *how* in the question refer to a particular outcome or endpoint, or perhaps instead a governing process, a set of rules or parameters to which the allocation must conform? Many approaches [12] focus on fairness and efficiency as global objectives. Furthermore, what are the objectives of the resource providing nodes and users? Are the providing nodes' objectives aligned, and do they align with the objective for the behaviour of the system as a whole, if one exists? If there is a conflict or tradeoff in achieving the objectives, how are these to be resolved?

In order to gain some perspective on these issues, it is interesting to consider Foster and Kesselman's [5] characterisation of computer systems as they scale. They note that in simple single *end machine* systems, resource allocation is typically dealt with at the operating system level, by a kernel or similar program which has absolute control over the resources in the machine. This enables it to achieve a tightly integrated system, but also provides a bottleneck, as resource requests must be fed through the kernel in order to be assigned. In *clusters*, many individual machines can communicate through message passing and file systems. Here increased scale is obtained at the expense of integration, as homogeneous nodes are controlled by a single machine responsible for job allocation. Larger still, *intranets* are characterised more by heterogeneity of nodes, which may be under administrative control of separate entities. Nodes may have different policies for use of their resources, different external demands and different capabilities. Here issues exist with regard to the availability of global knowledge. Nodes may attempt to map out the computing environment in order to plan the best use of resources, though the size and dynamic nature of such networks means that any one node is unlikely to have an accurate view of the system's current state [5]. The final category considered by Foster and Kesselman is perhaps the most interesting, that of *internets.* These forms of network span many organisations, locations and platforms and are large and heterogeneous. Here there is no central control and often no global objective with regard to resource allocation.

Nevertheless, the Internet provides the most potential when considering future approaches to massively scalable computing systems [12] and indeed the Folding@Home project [17] is an early example of this. However, Lai [12] argues that in achieving the goal of harnessing the scale of the Internet in order to provide accessible, generic computing platforms on a vast scale, resource allocation remains a problem.

2.1.1 Centralised Approaches

Classically, resource allocation objectives are achieved in a centralised manner, often relying on a single node responsible for, say, load balancing [18]. A balanced load, though by no means the only interesting outcome, can be used as an example of a desired resource allocation, an objective against which a particular approach to resource allocation may be tested. Load balancing is additionally in itself interesting, since it is useful in numerous real world scenarios, including telecommunications networks, road networks and electricity and water distribution networks. In many of these domains, even in very large scale systems, centralisation is the usual approach taken [18].

Resource allocation techniques can be divided into two groups, stateless and state-based [19]. Perhaps the most widely known and easily understood stateless approach, used to balance the load on web servers, is round-robin DNS. A more complex example is proportional share scheduling [12], in which resources are allocated to jobs according to a set of pre-determined weights. However, stateless approaches such as this are unable to take account of current server load or availability, leading to no guarantee that the desired outcome is achieved. Simple state-based extensions permit the usage of information about the resources being managed, and enable the proximity to the desired allocation to be measured. Examples of state-based resource allocation approaches include those which make use of geographical information and previous usage levels in order to determine an appropriate allocation of resource. A useful review and comparison of these approaches in the web server domain may be found in [19].

2.1.2 Problems with Centralised Approaches

Centralised resource allocation methods do however have a number of drawbacks [20]. These include:

- the requirement that the environment remain static while the central coordinator is calculating the optimal resource allocation,
- that the coordinator has global knowledge of the system and all nodes within it,
- that all coordination messages must route through the central point, counteracting the benefit from having resources distributed about the network, reducing scalability [21] and creating a fundamentally brittle system [22].

The Internet in particular is a dynamic network, where the first two requirements are highly unlikely to be met [5]. Brittleness may be mitigated against to a certain degree, by introducing backup coordinator nodes, however even in these cases the wider system is reliant upon the existence and performance of a small number of key nodes. Failure at these key points in the network may well cripple wider functionality, at best [23].

2.1.3 Decentralised Approaches

These drawbacks lead to the need for a truly decentralised approach to the allocation of resources that does not rely on a central coordinator [20]. In the field of grid computing, examples include Cao et al.'s [24] hierarchical approach, and TURBO [21]. In the latter, allocations are achieved through the reliance on altruistic behaviour between cooperating peers, which collaborate in order to reach a global objective.

Balanced overlay networks [25] are another effective and generic technique for balancing a load across a decentralised network. In this approach, resource providing nodes present an estimation of their availability to other local nodes to which they are connected. Newly arriving jobs take a random walk through the network and select the providing node with the highest availability. Upon accepting and completing a job, a provider node updates its availability estimate.

In decentralised peer-to-peer storage systems Surana et al.'s [26] approach may also be used. Here the case is considered when moving loads around the network also uses bandwidth. Their objective is therefore a balance between achieving an even load and minimising the amount of load moved. Their fully decentralised approach is, in effect, tantamount to performing a centralised calculation at each node, periodically requiring cooperative reassignment of a load, based on global knowledge of the system.

2.1.4 Cooperation, Non-Cooperation and Self-Interest

Critically however, previous decentralised approaches either rely on nodes' having complete global knowledge, or else cooperating to some extent in order to reach a shared objective [27]. As an example of this, in balanced overlay networks [25] resource users are self-interested within the bounds of the providers observed within their random walk, though the providers themselves are relied upon both to provide an honest and accurate account of their availability and to facilitate the random walk by exposing their local connections. In the case where such cooperation may not be relied upon, it is likely that the system's performance would deteriorate significantly. Similarly, Surana et al.'s [26] approach assumes both cooperation between nodes and global knowledge of the system.

A non-cooperative, decentralised approach to resource allocation does exist in the domain of downloading replicated files. Dynamic parallel access schemes [28, 29] make use of self-interested smart clients to increase the speed of file downloads. It is not clear however, how this approach might be generalised to other service-based systems.

Buyya et al. [27] argue that we may not always be able to rely on cooperation between nodes, for several reasons. Amongst these are the possibility that a node behaves erroneously, perhaps due to a software or hardware error such as a virus, unforeseen circumstances or an external fault. Large systems are also likely to be noisy systems, as data is lost or corrupted in transit and the likelihood of measurements being inaccurate or misreported increases. Finally, limits on and delays in information transmission mean that nodes' actions may be misguided or insufficient. Crucially, Khan and Ahmad [30] show that in any decentralised cooperative approach, global optima can only be achieved when all the nodes cooperate. It is for these reasons that this thesis looks towards an approach which does not rely on the cooperation of nodes.

Some confusion does exist within the literature however in the treatment of the terms *non-cooperative* and *self-interested.* It is important to note that non-cooperation does not imply self-interest. Indeed, in Khan and Ahmad's [30] study of various games-based resource allocation methods, they describe a model in which non-cooperative agents bid for jobs based on an honest estimation of the estimated time to complete a job. Their agents, though not cooperating, act without consideration of the benefit they expect to derive from their actions. Clearly, such a consideration is a prerequisite for self-interested behaviour and hence the behaviour they describe is not self-interested.

Indeed, it is the assumption that an agent will behave either cooperatively or non-cooperatively, regardless of its predicament, that is at odds with self-interest. A self-interested agent may behave either cooperatively or non-cooperatively at certain times. The key factor is that this decision will be made by the agent, based on whether it is in its own perceived interest to do so. In making this decision, the agent must therefore consider the benefit it expects to gain from the options with which it is faced. If it does not, it cannot be said to be truly self-interested.

Therefore, when considering systems where nodes are owned or administered by separate parties, such as the very large distributed systems discussed by Foster and Kesselman [5], rather than consider agents on a cooperative / non-cooperative spectrum, it may instead be more useful to know whether or not an agent is self-interested. If it is possible to assume this of nodes, then as will be discussed in the following sections, the models and tools of economics allow for a great deal of progress to be made.

2.2 Related Economics

When selecting components with which to compose an application in a service oriented architecture, or when deciding where to send an image processing task on a computational grid, appropriate resources may be available from a number of providing nodes. Similarly, large numbers of users may find themselves competing for access to the best resources, or a resource at a time more suited to their needs. If individual users and providers are acting in a self-interested manner, then the resulting interactions may be thought of as being an economy [13].

Indeed, large computer networks such as the Internet, made up of heterogeneous individuals with independent objectives can quite rightly be viewed as social networks as well as purely digital ones. It is perhaps of little surprise then that a social science such as economics might be useful in solving a problem such as decentralised computational resource allocation, since economics itself is concerned with the allocation of resources between individuals with different objectives in human societies. Therefore, in computational networks that are social, to what extent can economic theory be called upon in order to predict, and hopefully design the resource allocation behaviour of complex computational systems, where individual nodes are self-interested?

It is perhaps useful at this stage to present and define some relevant terminology, which is used throughout this thesis. Firstly, according to Begg et al. [31] economics is "how [a] society resolves the problem of scarcity" (p3). Furthermore, they state that "a resource is scarce if the demand at a zero price would exceed the available supply" (p5). This is exactly the scenario with which we are faced in the computational resource allocation problem. How then do humans attempt to solve it?

There have of course been a number of different approaches to this problem in human history, but one which is particularly dominant is the use of markets. Rothbard [32] describes a *free market* as "an array of exchanges that take place in society. Each exchange is undertaken as a voluntary agreement between two people or between groups of people represented by agents". Similarly, Begg et al. [31] define a market as "a set of arrangements by which buyers and sellers are in contact to exchange goods or services" (p32). The important factors here are that there is an exchange between two or more individuals, and that this exchange is voluntarily entered into by all participants.

In order to facilitate such exchanges, a particular type of good is often agreed to serve as currency, in which case the individual giving away currency in order to obtain another good is termed the *buyer*, while that which receives the currency and gives away the other good is termed the *seller*. It is of course not required that this formal delineation be present, though it has been argued [33] that an economy will evolve towards common agreement on a particular good to treat as currency, typically that which the individuals find easiest to retain and exchange widely without additional cost.

A mechanism through which voluntary exchanges between individuals are facilitated is called an *auction*, and though the rules for these can also emerge naturally in economies, the earliest theoretical description of an auction is given by Walras [34]. This Walrasian auction posits an auctioneer with full access to the participants' demand for a good at each possible price. The auctioneer is then able to calculate a single optimal price at which the exchanges take place, such that the quantity of the good supplied exactly matches the quantity demanded.

Mathematical micro-economic theory treats an economic system as being akin to a physical system, possessing a current state and forces acting upon it [34, 35]. Such a system has one or more equilibrium points, towards which in the absence of restrictive forces, the system will drift. At these equilibria, the system is at rest. Smith [36] first postulated that what he called an *invisible hand* would drive an unrestricted market economy towards greater prosperity, that free and voluntary actions of self-interested individuals leads to greater overall well-being.

This argument was refined by Walras and Pareto, who proposed a century later that through the perfect Walrasian auction, an economy will reach a Pareto-efficient equilibrium point [34]. At this stage, no individual may be made better off without making at least one other individual worse off. In reality however, due to either imposed or unavoidable barriers to free exchange, additional costs associated with performing a transaction, an imperfection in the ability of the auctioneer, or simply the cost associated with the time an auction may take, such perfect Walrasian auctions are often an unrealistic theoretical ideal [35].

A number of more realistic auction mechanisms exist in the real world, including the common English auction, found amongst other places on Ebay [37]; the Dutch auction; Vickrey auction and Continuous Double Auction, often used in financial markets. Cliff [38] gives a useful introduction to and critique of several auction mechanisms, including those listed here, while Friedman and Rust [39] provide a more detailed look at the Continuous Double Auction.

Purely electronic markets also make use of a range of auction mechanisms. In designing a mechanism, the aim is typically to achieve an efficient system overall, by making use of the self-interested nature of individuals. This is demonstrated by Phelps et al. [40], Byde [41] and David et al. [42] amongst others. For many, the ultimate aim of such research is the automation of the mechanism's design, appropriate to individual scenarios [43, 44, 45, 46]. Taking Cliff's [46] work as an example of this, a parametrised mechanism design space is specified, which may be searched in order to find high performing mechanisms for specific scenarios. Results from an evolutionary search demonstrate that classic, human-designed mechanisms are often far from optimal.

Importantly, though this thesis most certainly does not take a view on the applicability or ethics of Rothbard's [32] views or conclusions for human society at large, the Austrian School to which he belonged, along with the earlier Lausanne school, have long advanced the tradition of a mathematical, logical approach to economics, which is useful in this particular domain. In particular, the methodology of determining the outcome of economic activity as the product of interacting self-interested agents provides us with a highly useful tool. Nevertheless, it is important to stress that the application of economic theory to the design and analysis of computational systems has little to do with human economies, and mechanisms which work well in the engineering domain cannot be assumed to transfer value free to human society.

2.2.1 Utility and Rationality

One of the key ideas in economic theory is that of *utility*, the concept of relative satisfaction [34]. An individual, when faced with a choice between two or more alternatives, may assign each alternative a utility, allowing the alternatives to be compared with respect to each other. This way, economic decisions may be made and a self-interested individual may be described as being *utility maximising*, a term used often synonymously with economic *rationality*. This gives us the assumption that a rational self-interested individual will always make choices which maximise their own utility.

In human economics, the assumption of perfect rational utility maximising behaviour is a blunt instrument, and Simon [47] amongst others claims that in reality psychology, cognitive capacity and other factors produced by evolution will come in to play. This led to the development of the theory of *bounded rationality*, the idea that individuals act to maximise their individual utility within some constraints. These might be a lack of capacity to calculate the best action to take at a given time, a lack of full information concerning all the options available or any number of other factors. Indeed, Simon questions [47] whether full economic rationality would be indeed be evolutionarily optimal at all.

On a macro-economic scale this criticism is sometimes retorted by the *as if* argument popularised by Friedman [48]. This argument contends that though humans' decision making may well be bounded, random or indeed on occasions appearing purposefully economically irrational, when viewed on mass, humans behave as if they are indeed utility maximising. This debate continues amongst economists, however for the purposes of the design of economics inspired computational systems, it is merely helpful to be aware of the arguments surrounding individual economic decision making, especially as we will see later, with regard to the impact of different behaviours on global outcomes.

However, utility is by its nature subjective. Indeed, the concept of a scalar value for utility, the *util* was originally proposed as a unit of psychological measurement to enable the comparison of utility between individuals [35]. This approach was quickly discarded, since though it makes sense to consider whether Alice likes apples more than she likes bananas, it is meaningless to conceive of whether she likes apples more than Bob likes bananas. In modern economic theory, utility is more commonly thought of as an expression of relative preference by an individual between alternatives. These preferences may be denoted by means of ordinal or cardinal utility, the former of which records merely the order of preference between alternatives, while the latter attempts to place comparable numerical values on the alternatives. Care must be taken however when using cardinal utility, since the temptation can easily arise to assume that operations on cardinal utility values imply more than they are able to meaningfully convey. In practice, this means that the values represented by cardinal utility preserve preference orderings up to positive linear transformations [34]. Though cardinal utility does not provide a means of comparing satisfaction between individuals, it can be useful in individual agents' decision making processes.

Indeed, elements of decision theory build upon cardinal utility in order reason about tradeoffs between relative preferences and values. Keeney and Raiffa [49] in particular argue that utility models can be constructed to represent humans' preferences, and that these can be useful in aiding decisions in complex practical situations. Of course, many decisions require the consideration of numerous objectives, costs and the tradeoffs between them, and this is where approaches they advocate such as multi-attribute utility theory are particularly useful.

2.2.2 Relevant Game Theory

As with any discipline, economics has its tools of the trade, and one of the most important and useful of such tools in micro-economics is game theory. In essence, a game is what occurs whenever two or more rational individuals interact, and game theory is concerned with how the outcomes of these interactions relate to the individuals' preferences and the structure of the game [50]. Game theory was formally defined when Von Neumann and Morgenstern [51] published their seminal book in 1944, and subsequently developed throughout the 1950's by many others, including most notably Nash [52].

However, examples of what became known as game theory have existed in specific instances throughout history. Perhaps the most important of these were the early models of economic competition due to Cournot and Bertrand [53]. As far back as 1883, Cournot proposed that firms[1] producing equivalent goods, each knowing what price they can obtain for their good in the market, compete by making decisions about the quantity of the good to produce. Crucially, Cournot considered that firms act independently of each other, act strategically and seek to maximise their profit when making their decisions about what quantity of the good to produce. This is perhaps one of the first examples of analysis of the interactions of self-interested strategic individuals and resultant effects.

Bertrand [53] responded to Cournot by noting that the law of one price does not always hold, and therefore one firm may indeed lower the price it charges for the good and in doing so capture the entire market. By increasing the quantity it sells, a firm's profit would be sure to increase also. However, Bertrand further observed that a competing firm would, when acting rationally, respond to this by reducing the price it charged in turn. In doing so, it would expect to capture the entire market and increase its payoff. The other firm responds in kind, and so a price war begins, as firms compete to supply the entire market, by undercutting the other's price in turn. The outcome of this, predicted by Bertrand, would be that both sellers reduce their prices as much as they can bear. Assuming that the firms have equivalent costs, then they would reach an equilibrium when they were both charging their cost price. At this stage, neither seller can increase its payoff by unilaterally changing its decision. This is the Bertrand equilibrium [34], and complements an equilibrium discovered through similar analysis by Cournot in his model.

The framework provided by Von Neumann and Morgenstern [51] allows competitive games of this type to be described and reasoned about in a principled way. They claim that a game can be completely described in terms of the following pieces of information:

- The players of the game;
- For every player, every opportunity they have to move (make a decision);
- What each player can do at each move (their options);

[1] In economic literature sellers are sometimes referred to as firms.

- What each player knows at each move; and
- The payoffs received by each player for every combination of moves.

Each player is therefore faced with one or more decisions, and given the payoff structure of the game and the knowledge available at each decision point, is assumed to act such as to maximise its payoff. The actions a particular player takes at each decision point are defined by the player's *strategy*. A strategy can then be thought of as a set of rules which defines the player's actions, in any and all circumstances which may arise during the game. A strategy therefore defines a player's entire behaviour.

This is perhaps best illustrated by the canonical example of a game, the *Prisoner's Dilemma* [54]. In this idealised game, two individuals are in police custody on suspicion of having committed a crime and are separately questioned. Each is assumed to value their own freedom beyond that of their accomplice, and is therefore faced with a dilemma. Each player can either *cooperate* with the other by keeping silent, or else *defect* by providing a statement that the other player is guilty. Here we have the players of the game and the opportunities they have to move. In this particular game, both players also know nothing about the other's action, since their decisions are made simultaneously in isolation.

The prosecutor, who sets the structure of the game, provides the players with information concerning their payoffs for each action. These are best presented through a payoff matrix, an example of which is shown in figure 2.1. Here, the numbers represent the number of years in prison the player will be sentenced to, in each case. The values are negative, since strictly this is a payoff matrix, and the payoff associated with increasing time in prison is certainly increasingly negative. The payoff values are shown for the row player, here player 1. Since the game is symmetric, the same is true in reverse, and only one value needs to be shown in each space in the table. In this example, if both players were to defect and provide evidence to convict their accomplice, then each would receive five years in prison, since they would be found guilty, but had also pleaded guilty to the crime. If both players were to cooperate by staying silent, then each would receive one year in prison on minor charges only. However, if player 1 were to defect, but player 2 cooperated, then player 1 is allowed to go free for assisting

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	-1	-10
	Defect	0	-5

Figure 2.1: Sample payoff matrix for the Prisoner's Dilemma.

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	R	S
	Defect	T	P

Figure 2.2: Generalised payoff matrix for the Prisoner's Dilemma.

the authorities in their investigations, but player 2 is punished with ten years in prison, since he is both found guilty and refused to confess. The reverse is true if player 1 cooperates but player 2 defects.

More generally, the payoffs for a game to be considered to be the Prisoner's Dilemma are as given in figure 2.2, where R is the reward for cooperation, P is the punishment for defection, T is the temptation to defect, and S is the sucker's payoff, if the player cooperates and the other defects. These must be structured such that $T > R > P > S$.

In reasoning through a player's actions, it is clear that regardless of the opponent's choice, the rational choice is always to defect, since $T > R$ and $P > S$. Since this is true for both players, both will defect and neither could have done better by doing otherwise. This illustrates the concept of a *Nash equilibrium*, a strategy profile from which neither player may unilaterally deviate in order to increase his payoff [52]. In this particular case, defection is also a dominant strategy, since whatever the opponent were to do, it always remains the rational choice.

This is clear, however what is less so is the behaviour of rational players when the game is repeated [54, 55]. In the repeated case n rounds of the game are played, where $n > 1$ and an additional constraint is added to the payoffs, that $2R > T + S$. Each player is able to remember the actions of his opponent between rounds, and a player's payoff is the sum of the payoffs gained in each round. Here, due to backwards induction, the Nash equilibrium is still for each player to defect in every round [54], though since $nP < nR$ this does not give rise to the highest achievable payoff for either player.

Since the Iterated Prisoner's Dilemma was first proposed by Axelrod [54], numerous contributions have been made to understanding the dynamics of the game in evolutionary models, on networks and with greater numbers of players and choices [55]. Key to players' obtaining a higher payoff is the achievement of a level of cooperation between them; their developing an understanding with their opponent that they will not decide to defect but instead continue to cooperate. The players forgo the higher short term temptation T, which would likely end up in both receiving P, in exchange for a higher longer term payoff, ideally nR.

The Iterated Prisoner's Dilemma is relevant in economic games, particularly Bertrand competition, since it shows how cooperation or collusion between self-interested individuals can bring about a higher payoff for each. Recall that the Bertrand equilibrium, which is also a Nash equilibrium [34] occurs when self-interested sellers undercut each other's price in order to obtain a higher short term payoff, leading to both sellers having no option but to sell at their cost price. This is analogous to defection in the Prisoner's Dilemma game. Cooperating sellers could instead choose to maintain higher prices, if they could be assured of the cooperation of their competitors in doing the same.

In the majority of advanced market economies, such price-fixing collusion is considered illegal. Nevertheless, in oligopolies, markets where the number of sellers competing to provide a certain good is small, collusion may emerge implicitly simply by firms observing and responding to each other's actions. Cheung et al. [56] demonstrate this effect, showing how such implicit collusion may emerge naturally. Nevertheless, in markets with larger numbers of sellers, this becomes increasingly less likely to occur.

The possibility of collusion between sellers allowing a market to remain away from a competitive equilibrium is then a contribution that game theory makes to our understanding of economic models. In larger markets, as discussed by Cheung et al. [56] this can be assumed to not occur, but nevertheless it should be borne in mind especially in markets with smaller numbers of sellers. Most importantly however, game theory allows us to reason about the interactions of self-interested individuals, and in doing so identify and predict outcomes for a given economic system.

2.2.3 Agent Based Economic Modelling

An individual which performs an action in a particular situation is called an *agent*. In economics, the term agent is used specifically to refer to an actor who makes individual decisions within a model [34]. These decisions are typically expected to seek to maximise the agent's utility, within the constraints of their knowledge and ability. In this way, a rational economic agent can be seen as being equivalent to a player in game theory. As has been discussed, micro-economics in general and game theory in particular seeks to reason about the interactions of such agents and therefore find and predict outcomes in a given economic situation. However, game theory does not scale well. Indeed, considering the combinations of agents' actions and interactions in all but the smallest games quickly becomes exponentially hard. In larger systems, macro-economic models have classically been dominant, since they account not for individual agents, but instead attempt to model their aggregate behaviour [34].

However, an alternative to this which permits us to keep the interactions between individual economic actors in consideration is to use computational agent-based simulation to build generative models of economies from the ground up [57]. This approach is also applied with success in other social sciences [58, 59]. Both in facilitating this and in acting in real scenarios where situated agents act on behalf of humans, a further use of agent terminology, the *software agent* has come into being. Wooldridge and Jennings [60] describe a software agent as a being computer system in possession of (at least) autonomy, social ability, reactivity and pro-activeness. More precisely, they specify an agent as being:

- Clearly identifiable problem solving entities with well-defined boundaries and interfaces;
- Situated (embedded) in a particular environment - they receive inputs related to the state of their environment through sensors and they act on the environment through

effectors;

- Designed to fulfil a specific purpose - they have particular objectives (goals) to achieve;
- Autonomous - they have control both over their internal state and over their own behaviour;
- Capable of exhibiting flexible problem solving behaviour in pursuit of their design objectives - they need to be both reactive (able to respond in a timely fashion to changes that occur in their environment) and proactive (able to act in anticipation of future goals).

Sengupta et al. [61] argue that traditional economic analysis has indeed not been very successful in dealing with the complexity of markets with heterogeneous individuals. They claim that agent-based models are much better suited to such complexity, and in addition allow the modeller to focus on individuals' behaviour, for example using psychological models of consumer choice rather than purely statistical ones.

Multi-agent systems have also become a dominant paradigm for the design of distributed systems [43], where individual components fulfil their own independent goals. Such systems are often characterised by the conflicting objectives of agents, which must be resolved in order to find agreement. As we have seen, markets provide a mechanism by which individuals may settle on a particular agreement, in this case an allocation of resources.

Software agents are therefore useful for two related reasons. Firstly, they may be assembled into generative models of social and economic systems, through which results akin to those obtained game theoretically may be obtained. And secondly, they may be embodied in the software systems being modelled, as autonomous individuals, capable of acting in order to further their goals. In the specific economic case, both sellers and buyers may be represented in the a market by autonomous software agents.

2.2.4 Evolutionary Approaches to Economic Games

As was discussed in section 2.2.2 above, a player's behaviour in a game can be completely specified by its *strategy* [50]. This defines the player's behaviour in every and all possible situations in which it may find itself. In simple single round games such as the Prisoner's Dilemma, a strategy may be represented by a single bit of information which determines whether the player cooperates or defects. However, in even moderately sized repeated games and certainly in large ones, the strategy space, the domain which defines all possible strategies, becomes unmanageably large.

Consider for example the Iterated Prisoner's Dilemma. Here a player has two choices in each round. Therefore a player who takes no account of the actions of his opponent and continues regardless might represent his strategy in terms of n bits, where n is the number of rounds. A player who takes into account the previous round when making his decision would instead have to consider what his action would be in $2^2 = 4$ possible states. Taking account the previous r rounds, the number of states is 4^r [62]. If a player's strategy consists of a full set of rules, one for each of these possible states, then there are a possible 2^{4^r} strategies [62]. The size of the strategy space makes strategy acquisition for the Iterated Prisoner's Dilemma an interesting problem. Furthermore, since the iterated version of the game has no strategy which dominates all others [63], competing players are not faced with a static search problem. This problem has lead to a large amount of literature, much of which was motivated by computer based Iterated Prisoner's Dilemma competitions [63].

One approach which has proved effective in tackling the problem of learning a high performing strategy is evolutionary computation [55]. Evolutionary algorithms provide us with the tools to be able to search for solutions to complex problems such as this, and including those faced by economic agents. In the years since evolutionary computation was first pioneered [64, 65, 66], a number of techniques have now matured, and have been applied to a wide range of computational problems, including search, learning, optimisation, design and game theoretic analysis [67, 68, 69]. As we might therefore expect, there is more than one approach which could be employed here.

In strategy acquisition, co-evolutionary learning has proved particularly effective [15, 56, 59, 70] and theoretical work now exists which begins to quantify the expected generalised performance of learnt strategies against unknown opponents [71, 72]. Co-evolution is an extension of evolutionary computation in which the fitness function is in part dependent upon the actions of others [68]. Co-evolution has been used to learn strategies in repeated games such as the Iterated Prisoner's Dilemma [63] and also to find Nash Equilibria in games with continuous strategy spaces [73].

Co-evolutionary learning algorithms can be divided into two key families, the *Michigan* and *Pittsburgh* (or simply *Pitt*) approaches [74]. In the Michigan approach, the individuals in the population consist of separate rules, which together form the complete strategy. Conversely, in the Pitt approach each individual in the population is itself a complete strategy. Since rules are only tested as part of a whole strategy, the Pitt approach provides a convenient mechanism for assessing strategies when compared with each other, whereas the Michigan approach results in a difficult problem of payoff attribution. For this reason, Pitt-based algorithms have been the more widely used, and are the most easily understood.

Pittsburgh co-evolutionary learning has much in common with *evolutionary game theory*, a further tool used to analyse games. Though developed originally for use in biology [75], this approach is now also used in economic scenarios [44, 40, 76, 77]. By allowing a population of strategies to evolve, rewarding the most successful strategies by replicating them (potentially with some mutation), while culling the least fit strategies, evolutionary dominance and the stability of certain strategy profiles may be explored. Vytelingum et al. [77] provide a useful introduction to the applicability of evolutionary game theory to the analysis of agent-based markets. They detail how populations of buyers and sellers can evolve towards the dominance of certain strategies interacting through a continuous double auction. Kephart et al's [9] analysis of pricing behaviour uses a similar approach to analysing the dynamics of a simple market.

However, the issue of exponential growth of the strategy search space is fundamental. Even in large static markets, agents are faced with the problem of finding a good strategy from an immensely large search space. Furthermore, the generalisation issue [71, 72] illustrates that finding strategies which work well against previously unseen competitors is harder still. If agents are dealing additionally with the possibility that the market may be dynamic, as others arrive and leave or update their preferences, then extending a rule-based approach such as is used in the Iterated Prisoner's Dilemma [62] or in the smaller pricing model by Cheung et al. [56] becomes unfeasible.

A promising alternative to static evolutionary strategy search, with a greater focus on adaptability, is the treatment of action selection as a dynamic optimisation problem. Though developed as a technique for use on stationary problems, evolutionary computation has in recent years been used in dynamic environments [78]. In this case, the goal of the algorithm is no longer to find the optimum point in the search space, but to track its movement over time. Branke [78] provides a comprehensive introduction to techniques which have been suggested for the modification of traditional evolutionary algorithms to problems with dynamic fitness landscapes. The majority of approaches rely on the technique of retaining a certain amount of information about the landscape, through maintaining diversity in the population. A number of techniques have been proposed to ensure this, such as the periodic introduction of *random immigrants* to the population [79] and *triggered hypermutation* [80]. However, the random immigrants approach is a somewhat blunt instrument, which may disrupt the effectiveness of the search and prevent the algorithm from finding the global optimum [81]. Triggered hypermutation may be highly effective, but relies on certain domain-specific knowledge, such as what should be considered a trigger. A third option, when some pattern exists in the nature of the environment's dynamics, is the retention of knowledge of the landscape in the individuals themselves: a form of memory [81, 82]. Regardless of the specific approach taken, by not attempting to learn a strategy which considers every possible scenario in which the agent finds itself, but instead acting to maximise payoff in the nearer term, the complexity of the search space and hence the difficulty of the strategy acquisition problem may be greatly reduced.

2.3 Economics-Inspired Resource Allocation

As has been discussed, large computational networks such as the Internet, consisting of interacting nodes with independent objectives can be viewed as an economy. As a result, economics has been both applied and used as inspiration for techniques to perform resource allocation in such computational settings [13].

The actions and decisions of resource providing and resource using nodes can be automated by the use of software agents interacting in a market. The aim of a buyer agent might be to secure the fastest and most reliable resource at the lowest cost for its user. Conversely, a seller agent might aim to maximise the revenue for the resource provider, or perhaps generate high levels of business. Whatever the business strategy of the resource provider, the selling agent will be competing with similar agents from other providers for the same resource users. Each agent will therefore have to employ its own strategy for success in the market.

This idea is not new. In a discussion of lessons learnt from experience with load management in giant-scale web services, Brewer [18] proposes the idea of incorporating, into a request for a resource, a notion of its value or cost. It is argued that this, along with the use of smart agents, would allow for responsive adaptation in the presence of changes to the network, as well as graceful degradation. Similarly, Gupta et al. [83] argue that in the provision of virtually zero cost per-use computational services, a mechanism involving pricing and user self-selection is preferable to the alternative of provider or regulator enforced limits: rationing.

The application of economic ideas to resource allocation problems in computational systems is approached in the field of market-based control, an introduction to which is provided by Clearwater et al. [13]. Using the terminology of Casavant and Kuhl's [84] taxonomy of scheduling in distributed computing systems, this is a family of distributed mechanisms for dynamic global resource allocation.

Typically, resource owning or providing nodes are represented by selling agents, and resource users or tasks are represented by buying agents. Buyers then attempt to purchase sufficient resource to satisfy their task or user's requirements from the set of available sellers.

Sellers charge an amount of either real or artificial money for the resource, determined by their strategy and dependent on factors such as the quantity or quality of the resource being provided. Since self-interested buyers can be expected to pay more for resources which they desire more, and self-interested sellers will charge what they can get away with in order to maximise their payoff, resources will tend to go to those who value them the most.

Fundamentally, these approaches attempt to harness the rational behaviour of self-interested agents, which interact in some market environment in order to achieve resource allocation without reference to a central authority. Relying upon the theories of micro-economics discussed in section 2.2, through such repeated exchanges between utility maximising individuals, efficient resource allocations may be achieved.

2.3.1 Centralised Market Mechanisms

As in human economies, agents in a market-based computational system may interact through any of a number of different mechanisms [38]. Common examples include English, Dutch and Vickrey auctions, in which an auctioneer facilitates the bidding and determines the allocation of resources. A range of market-based resource allocation and control applications are described by Clearwater et al. [13].

Where scarcity exists on both the seller and buyer sides, double auctions such as the Continuous Double Auction and Clearing House provide an alternative approach [39]. Research in the field of automated mechanism design also suggests that other less obvious auction mechanisms may lead to more efficient outcomes in certain circumstances [46, 85, 86].

However, both Cliff and Bruten [15] and Eymann et al. [20] note that due to the mechanisms employed, a large proportion of market-based control systems are not truly decentralised, since they rely on a centralised price fixing process rather than the participants between them determining prices. This is true of Wolski et al.'s [14] G-Commerce model, which relies upon a central market maker. Cliff and Bruten [15] argue that the presence of such a centralised process or component removes the primary advantage of using a market-based system: its robust, decentralised, self-organising properties.

2.3.2 Distributed Market Mechanisms

A number of distributed auction mechanisms have also been proposed [87, 88, 89], which do not rely on one central coordinating node. These approaches reduce the fragility associated with reliance upon a single point, provide more scalability and allow for dynamic composition of auctions. Typically, either the central auctioneer is replaced by a number of local ones, which may communicate through some secure means, or else the auctioneer role is fulfilled by a spare, disinterested node. Double auctions for example, though relying on a specialist to match bids and asks [40], may be decentralised by the presence of multiple specialists between which the participants may choose [90]. These techniques do reduce bottlenecks at certain points within the network and the removal of a single node cannot lead to system-wide failure. However, similarly to the replicated round-robin DNS approaches discussed in section 2.1.1 above, the system is still largely reliant on a small subset of its nodes.

However, it may be possible in systems such as this to scale up the number of auctioneers or specialists, in order to achieve a suitable degree of redundancy and decentralisation. This issue is an active area of research and worth investigating further, though intuition suggests that in an ideal world a system which relies upon a set of super-nodes can never provide the level of robustness of a system without such a need, even if the super-nodes were present in abundance. Approaches such as this also raise questions of motivation for those acting as super-nodes, as participation fees for example are set by auctioneers in most cases [90]. Therefore, if an approach exists without the need for such complexity, it should be preferred.

A further alternative is that individual provider nodes themselves host independent auctions for their resources. This approach is applied to computational resource allocation in Spawn [91]. Here, users' agents bid in sealed-bid auctions hosted by providers' agents, for their resources. In order to be effective, this requires a high level of strategic ability on the part of buyers, as they must decide in which auctions to participate. Of course, consumers may win multiple auctions, and questions then arise of how to handle these situations. Literature exists which explores the dilemma faced by buying agents bidding in multiple auctions, such as that by Gerding et al. [92, 93] though again this thesis attempts a less complex approach.

2.3.3 Bargaining

Cliff and Bruten [15] conclude from their critique that, rather than depend upon a central node such as an auctioneer, market mechanisms should instead rely on the ability of intelligent agents to bargain between themselves in order to arrive at acceptable prices. This approach is taken in the AVALANCHE [94], and CATNET [95, 96, 20, 22] systems. These take inspiration from Agent-based Computational Economics (ACE) [57], an agent-based modelling technique which attempts to replicate the dynamics of human markets with complex cognitive agents.

These approaches are those which attempt to replicate human markets the most faithfully, since they rely on highly developed strategies, as agents negotiate bilaterally in order to determine the provision of a resource. It is likely in this approach that the development and operation of such strategies will themselves require significant computational overhead. While these approaches are indeed effective and widely applicable, if a simpler alternative exists, it should be preferred where possible. An additional point of interest is that in the mechanism used in CATNET [20], resource providing nodes are relied upon to forward requests to neighbouring hosts, without any consideration of the effect of this on their own interests. This appears to be at odds with the self-interested nature of the agents. However, the study of bargaining agents is a topic of ongoing research [97, 98, 99] and has a relevance beyond that which is considered in this thesis.

2.3.4 Retail Markets and the Posted Offer Mechanism

Though they do not discuss them in detail, Cliff and Bruten [15] also briefly mention retail markets as an alternative to auctions and bilateral negotiation. The mechanism used in modern retail markets is usually referred to as the *posted price* or *posted offer* model [100, 101], though in online content delivery it is sometimes referred to as the *quoted price* model [102]. It is a fully decentralised approach to the determination of price without the need for complex bilateral negotiation, and provides a potentially simpler alternative.

Wang [103] provides an interesting comparison of auction-based and posted offer selling, and shows that auctions are more commonly used in human markets where there is a greater dispersal of valuations of the good amongst the buyers. Where buyer valuations are more similar however, he favours the posted offer market mechanism. This can be reconciled with the idea that according to the most common mechanism design objectives, there exists no single dominant mechanism [86]. For an example of this in a specific case, the impossibility result due to Myerson and Satterthwaite [104] shows that no double auction can simultaneously be efficient and budget balanced while also ensuring that at least one participant would not be better off using a different mechanism. It is therefore appropriate that research into computational resource allocation continues to consider the impact of a range of mechanisms.

Plott and Smith [100] describe the posted offer mechanism as being a process in which sellers of multiple units of a good each post one price or offer, and buyers subsequently respond by stating the quantity which they wish to purchase from each seller. Exchanges then occur between buyers and sellers at these price and quantity values. Technically, the reverse process in which buyers quote prices and sellers state quantities is also a posted offer mechanism, though is less commonly encountered. Importantly, price quotations cannot be changed during the exchange period.

Some examples of using similar mechanisms in computational resource allocation do exist in the literature, though they are not faithful implementations and make additional assumptions. Chavez et al. [105] use an approach of this type in Challenger, where offers are broadcast to the nodes in a network, though instead of using price, bids contain an honest reporting of a job's priority. However, this honesty means that there is no competition between nodes and as with Khan and Ahmad's [30] approach discussed in section 2.3 this not self-interested behaviour. Xiao et al. [106] describe their system GridIS, in which buyers broadcast job requests and sellers reply by posting offers to perform them at a price. However, the behaviour of the sellers used requires certain global information in determining their price, both in the form of the latest accepted market price, which in a posted offer mechanism is considered to be private information, and also the level of aggregate supply of all the providers in the network. Again, the assumption of private information forbids this too.

Perhaps the most faithful implementation of the posted offer mechanism in decentralised computational resource allocation is that by Kuwabara et al. [107], though they do not describe it as such. They propose an approach which in which sellers quote prices for their resources, and buyers subsequently decide the quantity (which may be zero) to purchase from each seller. Their analysis determines the quantities provided at the equilibria at which the markets arrive, and present this as a stable outcome allocation of resources. This is indeed fully decentralised, since no central component, such as an auctioneer or specialist is used; prices are determined privately by the sellers and then posted via a broadcast mechanism.

2.4 Adaptive Pricing

As in human retail markets, key to the efficiency of the posted offer market approach used by Kuwabara et al. [107] is the assumption of iterated transactions. For example, where a large number of buyers arrive over a period of time, each seller is able to adapt its price such that the payoff from its transactions is maximised. Sellers therefore compete on price, over time. From a seller's perspective, we are therefore faced with the question firstly of finding the best price, and secondly how this price should be adapted in response to market conditions. Of course, other sellers will also respond, as prices co-adapt. This is therefore a type of *adaptive pricing game*.

2.4.1 Approaches to Adaptive Pricing

Critically however, though the sellers in Kuwabara et al.'s [107] system are non-cooperative, they are not self-interested, since they do not consider their payoff when determining their price. Instead, a seller node's previous usage is encoded directly as the seller's next price, with no regard to the payoff it is likely to obtain by doing so.

Jagannathan and Almeroth [102] observe that in content delivery markets both online and offline, posted offer markets are frequently used. They also argue that pricing strategies from offline markets may not transfer well to online ones, and this agrees with Kephart et al.'s [9] argument that with the advent of automated software agent traders, market dynamics can be expected to be quite different from those involving humans. By considering a scenario where mirrored file servers operate from a number of locations around the world, Jagannathan and Almeroth [102] show that charging and adapting prices can be used to control the resource usage. For example, European servers may be busier and hence slower when North American servers are not, due to the time difference. By setting an appropriate pricing scheme, they show that this can be mitigated against.

The majority of the computational resource allocation models discussed in this chapter, with the exception of that by Surana et al. [26] assume that the cost of replication of data is negligible. Therefore, other than by taking into account fixed costs, in this idealised world payoff maximisation is equivalent to revenue maximisation. Reminiscent of Bertrand's [53] criticism of Cournot, Jagannathan and Almeroth [102] also note that on the Internet, sellers may quote different prices to different buyers, since each offered price is hidden. Indeed, this point also enforces that argument that in decentralised markets of this type, the last accepted price in the market will be private information.

Jagannathan and Almeroth [102] describe three families of pricing strategies:

Static Pricing In static pricing, the seller charges one price for all customers independent of time. This has the advantages of being simple and appearing fair, however finding the optimal price to charge is non-trivial. Additionally, this is highly likely to be suboptimal, since especially with heterogeneous buyers, there is a likelihood that at least one of them would have been prepared to pay more and hence increase the seller's payoff.

Discriminatory Pricing In this scheme a different price is quoted to each customer, based on the seller's best estimate of what they are willing and able to pay. In the short term this can maximise revenue, though in the long term, since this appears unfair, Jagannathan and Almeroth argue [108] that customers may be put off. They also claim that this is illegal in real markets in some jurisdictions, and as a result they are not keen on this approach.

Dynamic Pricing Here a price is quoted which changes over time, but does not vary by customer. Prices may change based on current system load, customer arrival rate, or other external factors. Dynamic pricing is particularly useful when information about the differences between individual customers is not available, or too expensive to store or process. This method also has the advantage of appearing more fair to customers, as they can all obtain the same price by adopting the same behaviour. This method can also be used to experiment with pricing. For example, by adopting a number of test prices at various times, the seller can explore the market in order to find a price which maximises revenue, given the current set of buyers.

Dynamic pricing behaviour can therefore be seen as strategic in the game theoretic sense. As was discussed in section 2.2.2, one effective way to computationally model multiple players' strategic behaviour, and hence self-interested dynamic pricing behaviour is with co-evolution. In the study of markets, co-evolution has indeed been used to optimise parametrised bidding and bargaining strategies [44, 94, 109, 110]. In co-evolutionary learning, this optimisation itself is dynamic, as it is performed against a moving optimum, as the competitors and potentially also buyers also update their strategies.

Price [111] demonstrates that, rather than optimising the parameters of a particular bidding strategy, certain classic competitive behaviour can be achieved by co-evolving prices directly. By drawing an analogy between payoff and evolutionary fitness, dynamic pricing by co-evolution is used to to drive competition in the market. In a sense, from a seller's perspective, evolution itself is the strategy. Amongst other examples, Bertrand competition is demonstrated by competitively co-evolving two sellers' prices.

Price's [111] approach appears initially highly attractive to system designers. The evolutionary algorithm employed by an agent does not require a model of the environment, since evolutionary fitness is instead obtained from the live market. It is a lightweight approach, which does not even require knowledge of the number or characteristics of competitors or customers. However, to date very little work has been done to investigate the suitability of this approach for computational resource allocation tasks.

2.4.2 Impacts of Adaptive Pricing

Some work analysing the effects of dynamic pricing in computational markets does exist in the domain of e-commerce. Greenwald and Kephart [112] explore a similar idea in the study of what they term *ShopBots*, which act for consumers in a model of an electronic media market, checking many sellers (online shops) and comparing prices. This enables customers to find the best deal. As a consequence of this, they also predict the advent of *PriceBots*, which act for sellers, automatically and dynamically setting prices to maximise their utility, as was described in the previous section. In this context, their work goes on to study the dynamics and predicted outcomes in such a system.

Their model considers two types of buyers. Firstly they consider buyers who will buy from any seller whose price is lower than the buyer's valuation of the good being offered, typically chosen at random. Secondly, they consider buyers who proactively seek out the seller making the best offer. They term these *type A* and *type B* buyers respectively, though in this thesis the more descriptive terms *time savers* and *bargain hunters* are preferred.

Their game theoretic analysis shows, as predicted by Bertrand's model, a population consisting of entirely type B hyperrational *bargain hunters* leads to an equilibrium where sellers charge their marginal cost price [112]. Conversely, they find that a population of only type A *time savers* allows sellers to behave monopolistically, charging a price equivalent to the buyers' valuation.

In Kephart et al.'s [9] earlier related work, two models of sellers' adaptive pricing behaviour are studied and Greenwald and Kephart [112] add a third. Firstly a rather unrealistic *game theoretic* behaviour is considered, in which sellers are able to take account of full knowledge of the buyers' and competing sellers' decision making processes. Secondly, a myopically optimal, or *myoptimal* strategy is used, where sellers have access to full knowledge about the current state of the other agents and their decision functions, but do not anticipate that other competing sellers will respond to their price changes or attempt to foresee the impact of their actions. *Myoptimal* sellers perform an analysis of the immediate consequence of each possible price they could charge, before making a decision. They are therefore able to behave fully rationally, or optimally, within the myopic constraint that they cannot consider the future. The third model is termed *derivative following*, which employs a simple hill climber pricing algorithm with profit as the objective function, and is therefore a somewhat simpler alternative to Price's [111] evolutionary pricing algorithm. Unlike the *game theoretic* and *myoptimal* methods, both of these *generate and test* approaches do not require knowledge of buyers' preferences or decision making processes, nor those of other sellers. This general approach is therefore likely to be much more representative of real scenarios where information is indeed private.

In the particular model studied, the market equilibrium reached depended also on the model of seller chosen [9, 112]. *Game theoretic* sellers, able to make use of full knowledge were unsurprisingly found to reach the Nash equilibrium. *Myoptimal* sellers were observed to engage in a price war, and in the case where the payoff landscape was multi-modal, the price war was found to repeat in an unending cycle [9]. *Derivative following* sellers produced an outcome similar to that observed by Cheung et al. [56], where implicit collusion allowed each to charge the monopolistic price.

Of particular interest is the observation by Kephart et al. [9] that *myoptimal* sellers in a multi-modal payoff landscape lead to a never ending series of price wars. These occur for a similar reason to that for defection in the Prisoner's Dilemma, that they may increase their short term payoff by undercutting each other on price. However, once the sellers' offered prices near their costs, rather than settle into equilibrium, it becomes rational for one of them to instead opt for a different section of the market, a different peak in the payoff landscape. They argue that the instability generated by this outcome is potentially a major obstacle for the success of electronic markets which rely on this kind of mechanism, since they foresee that multi-modal payoff landscapes will pervade real world embodiments.

2.5 Summary and Conclusions

In summary, emerging paradigms for the development and deployment of massively distributed computational systems allow resources to span many locations, organisations and platforms, connected through the Internet. In such systems, both resource providing and resource using nodes may arrive, organise and dissipate, as computational capabilities are formed and reformed as needed, without reference to a central authority or coordinator.

As these systems mature, it is predicted that the majority of their interactions will be carried out by autonomous software agents on behalf of their owners. In such distributed systems, where there exists a distribution of work to be done or resource to be provided about a network of nodes, neither control nor even full knowledge of key resources may be assumed, as they may be owned or administered by different organisations or individuals and as such have independent objectives. There is a need to find novel ways to understand and autonomically manage and control these large, decentralised and dynamic systems. As part of this, there remains the problem of how to allocate distributed resources amongst the nodes.

Economics provides a method for modelling such problems and reconciling conflicting nodes' objectives. In particular, game theory is a useful tool with which to reason about the interactions between self-interested agents. A number of different approaches to implementing this have been proposed, including single and double sided auctions. However, these typically either require a centralised price fixing process such as an auctioneer or specialist, or else regional super-nodes able to perform this function in a distributed manner. Both approaches require information to be channelled through one or more coordination points, raising questions of robustness and potentially creating bottlenecks. An alternative to this is bilateral bargaining, and this shows a great deal of promise as a fully decentralised approach, though it seems likely that this requires highly complex agent capabilities throughout the system, which will come with their own computational overhead. Furthermore, when agents are unable to fulfil this role, they will most likely be disadvantaged.

The simpler retail-inspired posted offer market mechanism provides a further promising alternative. Here agents are not required to possess complex strategic capabilities and indeed classic economic outcomes have been observed when sellers use evolutionary algorithms to decide their prices. The success of evolutionary computation in general and co-evolution in particular at strategy acquisition and learning in repeated games make it an ideal candidate for this, and evolutionary algorithms have been used in the design of market mechanisms and agent strategies alike.

However, a key characteristic of the systems being modelled is their sheer size. Though a strategy space may be defined and searched for small to medium sized economic games, as the system scales, the strategy space typically grows exponentially, making offline search for a good generalised strategy harder. There is therefore a need to investigate further the impact of more lightweight online evolutionary approaches to offer generation, and to assess the suitability of this approach, or lack thereof, for achieving resource allocation objectives.

Previous results are mixed, indicating that the approach may be able to generate outcomes predicted by classical economic theory, but also that effects such as price wars and collusion may alternatively induce unstable or undesirable outcomes instead. This thesis aims to contribute towards this understanding, comparing game theoretic results with those from evolutionary simulations in a variety of scenarios and with different resource allocation objectives. It is hoped that the insight gained from the results presented will be useful in both predicting outcomes from resource allocation systems embodying the models studied, and also provide a greater understanding of how market agents may be deployed in order to obtain these outcomes.

CHAPTER 3

A MODEL FOR DECENTRALISED COMPUTATIONAL RESOURCE ALLOCATION USING SELF-INTEREST

Essentially, all models are wrong, but some are useful.

George E. P. Box

In this chapter, the retail-inspired posted offer market mechanism is described and a model for computational resource allocation based on this mechanism is developed and analysed. Section 3.1 defines the posted offer mechanism itself, as used throughout this thesis, and formulates the resource allocation problem to be studied. In section 3.2 the model is analysed game theoretically for a small number of agents, and it is shown, with reference to Bertrand competition, how the approach may be used to achieve a balanced load in a computational setting. Section 3.3 concludes the chapter with a discussion, highlighting why some of the assumptions about the agents in the analysis may be unrealistic and raising important questions which emerge from this.

Throughout the analysis, a number of models of behaviour of resource using buyer agents are considered. Their effects on the qualitative outcomes of the model are compared and investigated. As such, this chapter provides both an introduction to the problem and a theoretical underpinning for the results presented later in the thesis.

3.1 Problem Formulation

3.1.1 Scenario and Objectives

As discussed in section 2.2.2, Bertrand's [34] model of economic competition is one of the simplest to account for the interactions between individual sellers who compete on price to provide homogeneous goods. The retail-inspired posted offer mechanism is qualitatively similar to Bertrand's model, in that it also accounts for sellers that compete on price to provide a homogeneous good to a population of buyers. In the types of computational resource allocation problems investigated here, the good is considered homogeneous, since the buyers do not care from whom they purchase a functionally equivalent resource[1].

However, unlike many other market-based resource allocation mechanisms, such as those discussed in section 2.3, this thesis focuses on how to achieve a particular outcome resource

[1]At this stage the reader may be concerned by the lack of account of non-functional attributes. These are considered in chapters 6 and 7.

allocation in a given scenario. The approach taken begins from the starting point of a desired allocation of resources which the system designer or owner wishes to achieve. An artificial market is then created in order to bring this allocation into effect, under the assumptions of decentralisation and self-interest.

Consider first a scenario consisting of a set of resource providing nodes, S, each member of which provides an equivalent, quantitatively divisible resource π, which may vary only in price. The members of S are assumed to be self-interested. Subsequently imagine a large population of resource users or buyers, B, each member of which aims to consume some of the resource π, at regular intervals.

If s_i is a node in S and b_j is a node in B, q_{ij} is used to denote the quantity of the resource π provided by s_i to b_j. The total quantity of π provided by s_i at a given instant, its ***load***, l_{s_i}, is therefore:

$$l_{s_i} = \sum_{j=1}^{|B|} q_{ij} \,. \tag{3.1}$$

The first resource allocation objective considered is a balanced load, such that at any instant, each resource providing node in S is providing an equal amount of π across the population of resource users. A particular resource allocation such as this, a configuration for the provision of π by the nodes in S at a given instant, may be expressed by the vector $\vec{L_S} = \langle l_{s_1}, l_{s_2}, ..., l_{s_n} \rangle$, where $n = |S|$. For convenience and ease of comparison between scenarios, this vector is often normalised by the total resource being provided. An evenly balanced load may therefore be written as $\langle \frac{1}{n}, \frac{1}{n}, ..., \frac{1}{n} \rangle$.

Though this is a trivial problem when central control or cooperation may be assumed, here the objective is to achieve this using only self-interest, in a fully decentralised manner with no central or regional control, and with only private information available.

It is worth noting at this stage that alternative approaches exist to the measurement of load. For example, we may wish to adopt a longer-term view, as in the *long-term proportional share* approach described by Lai [12]. In this alternative approach, load is averaged over a realistically set sliding time window. Casavant and Kuhl [84] also discuss the need to determine an appropriate measure of the load for the particular system. At this stage however, the instantaneous load as described above is used.

3.1.2 Mechanism and Assumptions

A posted offer mechanism is used to decide what quantity of the resource π is provided to which user node and from which provider node. At a given instant, a resource providing node, $s_i \in S$, advertises π at the price $p^{\pi}_{s_i}$ per unit via a broadcast mechanism. Each resource user, a buyer in this case, then has the option of purchasing some of the resource π, should it be in their interest to do so at the price offered. The system iterates, with sellers able to independently adapt their prices to the market conditions over time. This approach was first described in [1].

At this stage, the effect of the embodiment of the approach in any particular application or network environment is not considered, since this would make it unnecessarily difficult to analyse and understand the underlying behaviour of the model. With this in mind, a number of simplifying assumptions are taken:

1. That the system proceeds synchronously in discrete time-steps,
2. that each buyer desires exactly one unit of π per time-step,
3. that the actual provision of π may be regarded as instantaneous, such that it does not interfere with the mechanism,
4. that each seller has sufficient quantity of π available to satisfy all the buyers in B should it be so requested, and
5. that network connectivity is uniform.

The first two assumptions are present at this abstract stage in order to aid the analysis of the system, and though their presence has not been fully investigated, there appears to be no obvious intuition as to why they should alter the underlying behaviour being demonstrated.

Assumptions 3 and 4 may not be appropriate for certain embodiments of the approach, however are representative in theory of the provision of information-based services such as HTTP requests, and are present in other related work such as Kuwabara et al.'s [107] model discussed in chapter 2. Finally, assumption 5 replicates the network conditions found in Wolski et al.'s [14] G-Commerce system.

Each time-step, each buyer, if it chooses to buy, may purchase any amount of π from any number of resource providers in S, subject to the constraint that the total amount purchased per time-step is equal to exactly one unit (as per assumption 2). If no offer from any seller in S is acceptable, the buyer may instead purchase nothing. These constraints mean therefore that $\sum_{i=1}^{|S|} q_{ij} \in \{0, 1\}$ for all $b_j \in B$.

3.1.3 Buyer Behaviour

Both buyers and sellers accrue a payoff, or utility gain, from their interactions in the marketplace. For buyers, this is deemed to be the value they associate with the price paid subtracted from the value they associate with the purchased resource. If buyer b_j's unit valuation of π is denoted by $v^{\pi}_{b_j}$, then its payoff from a unit transaction with s_i will be $v^{\pi}_{b_j} - p^{\pi}_{s_i}$. Since any buyer accepting a price above $v^{\pi}_{b_j}$ would lead to a negative payoff, this is its reserve or limit price. From a buyer's perspective, if a seller's price would not lead to a negative payoff for the buyer, then the price is described as being ***acceptable***. S_{b_j} is used to denote the subset of S which contains exactly those sellers in S whose price is ***acceptable*** to buyer b_j. When buyers are homogeneous in so far as they have the same reserve prices, such that $v^{\pi}_{b_j} = v^{\pi}, \forall b_j \in B$, a set of sellers acceptable to the buyer population B exists, and is denoted as S_B. Of course $S_B \subseteq S$, or more precisely $S_B = \{s_i : s_i \in S,\ p^{\pi}_{s_i} \leq v^{\pi}\}$.

As with sellers, buyers are assumed to be self-interested and boundedly rational, at least insofar that they prefer higher payoffs to lower ones. As with real economic actors, this is manifested through the following of some strategy. Here, the strategy incorporates a decision function, which given a situation describes the quantity (which may be zero) to buy from each seller. This is a similar approach to that taken in by Greenwald and Kephart [112], in which buyers may be either hyperrational *bargain hunters*, seeking out the best possible price, or else *time savers* who will purchase from any *acceptable* seller, chosen at random[1]. The possibility of complex buyer decision functions means that there may not be a straightforward mapping between sellers' prices and buyer valuations, and the subsequent outcome allocation. Determining the outcome is therefore non-trivial.

Though buyers may adopt any of a number of behavioural strategies, throughout this thesis three representative buyer types are considered. These are Greenwald and Kephart's [112] hyperrational *bargain hunters* and *time savers* from and a further type, a risk-averse *spread buyer* behaviour [1]. These are now described.

Bargain Hunters

Bargain hunters always attempt to maximise their instantaneous payoff. In each iteration, they check the prices of all the sellers, selecting the one seller which provides the most attractive offer (i.e. the lowest price). If this price is acceptable, then the buyer purchases its entire unit of π from that seller. In the event that more than one seller provides an equally attractive and acceptable offer, the buyer purchases an even proportion of π from each such seller. This is the basic model of consumers used by Bertrand [34].

Time Savers

Time savers do not check the price of every seller in the system when deciding from whom to buy. Instead, they select a seller at random, and if its price is acceptable, then they purchase the entire unit of π from that seller. If it is not, then they continue selecting previously unchecked random sellers until they find an acceptable price. If no seller has an acceptable price, then they purchase nothing.

[1]Greenwald and Kephart [112] refer to *time savers* as *any seller* or *type A* buyers.

Spread Buyers

Spread buyers are rudimentarily risk-averse, preferring to spread their purchases across a number of sellers. At each time-step, the buyer looks at all the available offers, and purchases a proportion of π from each seller with a price below $v^{\pi}_{b_j}$, relative to the expected utility gain from purchasing from that seller. Specifically, the quantity purchased by buyer b_j from seller s_i is determined according to the following calculation:

$$q_{ij} = \frac{(v^{\pi}_{b_j} - p^{\pi}_{s_i})}{(nv^{\pi}_{b_j} - \sum_{k=1}^{n} p^{\pi}_{s_k})} \,. \tag{3.2}$$

Spread buyers only consider those sellers with an acceptable price.

Alternative Buyer Behaviours

It is worth reinforcing that although three buyer behaviours are considered here, many other potential behaviours will exist. For *spread buyers*, for example, an alternative might be to motivate risk-aversion through the model itself, by incorporating a real probability of sellers defaulting or behaving dishonestly. However in order to avoid complicating the model at this stage, the assumption of risk averse behaviour is used to gain extra clarity instead.

A behaviour with a high level of practical relevance might be that of *sticky buyers*. *Sticky buyers* prefer to continue using resource providers which they have previously used. They might switch to another seller if the long-term gain were greater than some threshold, and this could be motivated by an explicit cost to switching or discount for loyalty, amongst other things.

However, at this stage these are left as items for future research. Indeed, given the vast possibilities in agent behaviour, the analyses outlined in this thesis are intended to lay out a methodology for the determination of outcome allocations for a given set of agents. While some results are specific to example scenarios, it is the method outlined in these examples which forms a major contribution of the thesis.

3.1.4 Seller Behaviour

Sellers also receive a payoff, defined by their payoff function. Seller s_i's payoff is denoted as P_{s_i}. In its simplest form, this is its revenue from the sale of π:

$$P_{s_i} = \sum_{j=1}^{|B|} p^{\pi}_{s_i} q_{ij} \; , \tag{3.3}$$

or indeed

$$P_{s_i} = p^{\pi}_{s_i} \times l_{s_i} \; . \tag{3.4}$$

Clearly, a seller wishing to maximise its revenue would aim to increase both its price and the quantity of its resource sold to the buyers, its market share. However as we have seen from the buyers' behaviour, the market share will depend upon the relationship between its price and those of its competitors, specifically a higher price is likely to lead to a lower market share. The following section investigates this tradeoff.

3.2 Predicted Outcomes

One motivation for employing an artificial market is that competition between self-interested sellers drives the system towards equilibrium. It is at this equilibrium that the system is stable in the long term, and thus the allocation of resources in this stable state is referred to in this thesis as the outcome resource allocation. Of course, in realistic scenarios, continuing changes to the system may lead to the existence of attractors rather than equilibria in the classic sense, though this case is not considered at this stage.

The model described here is, in essence, a generalised version of the Bertrand game [34]. As discussed in chapter 2, the classic Bertrand game consists of two sellers, both of whom offer to sell a certain homogeneous good to a population of buyers. Each seller must decide what price to charge for the good, and then supply the quantity subsequently demanded by the buyers. The buyers in the classic Bertrand game behave hyperrationally, as with the *bargain hunters* studied here, always buying from the seller with the lowest price, or half from each seller if the prices are identical.

In this game either seller can take the entire market by offering a price only fractionally lower than its competitor. However, since this applies to both sellers, the non-cooperative Nash equilibrium for the game is for both sellers to charge as little as possible, their zero-profit price. If each seller's costs are equal, then the equilibrium price for each seller will also be equal. This leads to the sellers sharing the market equally at equilibrium, and it is this basic idea which provides us with a balanced load in the simplest case.

However, in the more general case, where buyers may follow any of a number of strategies, calculating the expected outcome resource allocation may be a more complex task. Here a methodology is followed of determining the sellers' best response at each iteration, by solving payoff equations constructed from the given buyer behaviour. This enables us to identify the Nash equilibrium outcome, where each and every seller's best response is equal to its previous position.

In the following discussions, it is assumed that the buyers do have an identical reserve price, $v^{\pi} = 300$, and therefore that we have a single acceptable set of sellers, S_B. Any seller in S but not in S_B will of course attract no buyers at all, and will hence receive no payoff and have a load of zero. For the sake of clarity, in the remainder of this section, only those sellers in S_B are considered.

3.2.1 Bargain Hunters

Let us first consider a scenario with two identical resource providing nodes, such that $S = \{s_1, s_2\}$, each with costs of zero. Recalling the sellers' payoff function, given in equation 3.4, we have that

$$P_{s_1} = p^{\pi}_{s_1} \times l_{s_1} \,. \tag{3.5}$$

and

$$P_{s_2} = p^{\pi}_{s_2} \times l_{s_2} \,. \tag{3.6}$$

As in Bertrand competition, B is a large population of hyperrational buyers, ***bargain hunters***, as described in section 3.1.3. Recalling the decision function for these buyers, and the assumption that each buyer wishes to purchase exactly one unit of π, we may therefore say that

$$P_{s_1} = \begin{cases} |B| \times p^{\pi}_{s_1} & \text{if } p_{s_1} < p_{s_2}; \\ 0.5 \times |B| \times p^{\pi}_{s_1} & \text{if } p_{s_1} = p_{s_2}; \\ 0 & \text{otherwise.} \end{cases} \tag{3.7}$$

and the equivalent for s_2 respectively.

From a game theoretic perspective, given an observed value for their competitor's price, both s_1 and s_2 will wish to respond with the best response. In this case, this will be to undercut the competitor's price, if possible, in order to receive the payoff given by the first case in equation 3.7. The competing seller will of course act similarly, leading to a price war where each undercuts the other until their zero-payoff price is reached. Assuming that a seller would rather not participate than receive a negative payoff, once $p_{s_1} = p_{s_2} = 0$, the rational course of action is to maintain a price of 0, accepting the second case.

Recalling that the current load on a resource providing node is given by equation 3.1 above, we therefore have that at equilibrium,

$$l_{s_1} = 0.5 \times |B| \,, \tag{3.8}$$

and

$$l_{s_2} = 0.5 \times |B| \,. \tag{3.9}$$

This is indeed an evenly balanced load, i.e.

$$\vec{L_S} = \langle \frac{1}{2}, \frac{1}{2} \rangle. \tag{3.10}$$

The theory of Bertrand competition (which is described more fully in [34]) demonstrates that when competing on price alone, two sellers are enough for the perfectly competitive outcome described here. Since the same logic applies to larger number of sellers, this evenly balanced outcome also holds for larger systems under the same assumptions. This idea was was first presented in [1] and elaborated upon in [3].

3.2.2 Time Savers

Intuitively, a population of *time savers* will posses less of the *all or nothing* nature of *bargain hunters*, as each will prefer potentially any seller whose price is acceptable. Considering the simple two node example described above, what outcome should we expect with a population of *time savers*? Recalling that only those sellers in S_B are considered at present, the payoff for s_1 and s_2 should be expected to be

$$P_{s_1} = \frac{p_{s_1}^{\pi}}{|S_B|} \tag{3.11}$$

$$P_{s_2} = \frac{p_{s_2}^{\pi}}{|S_B|} \tag{3.12}$$

Here, unlike with *bargain hunters*, there is no advantage for a seller in undercutting the price of a competing seller, since this will only serve to reduce its payoff. Instead, the dominant position is to charge the highest possible price while still remaining in S_B; the equilibrium is at $p_{s_1} = p_{s_2} = v^{\pi}$.

Similarly to *bargain hunters* however, since $p_{s_1} = p_{s_2}$, then $\vec{L_S} \approx \langle \frac{1}{2}, \frac{1}{2} \rangle$. Note that due to the probabilistic nature of the buyers' decision function, the allocation will tend towards this as the probabilities average out.

3.2.3 Spread Buyers

For a population of *spread buyers*, as described in section 3.1.3, the sellers' payoff functions for the simple two node case are

$$P_{s_1} = \sum_{j=1}^{|B|} \frac{v^\pi - p^\pi_{s_1}}{2v^\pi - (p^\pi_{s_1} + p^\pi_{s_2})} \times p^\pi_{s_1} \,, \tag{3.13}$$

and

$$P_{s_2} = \sum_{j=1}^{|B|} \frac{v^\pi - p^\pi_{s_2}}{2v^\pi - (p^\pi_{s_2} + p^\pi_{s_1})} \times p^\pi_{s_2} \,. \tag{3.14}$$

Sellers s_1 and s_2 will each then attempt to maximise their respective payoff function as before. The outcome resource allocation occurs when the system is at equilibrium. Figure 3.1*a* illustrates an example payoff function for s_1, when $v^\pi = 300$ and $p^\pi_{s_2} = 250$.

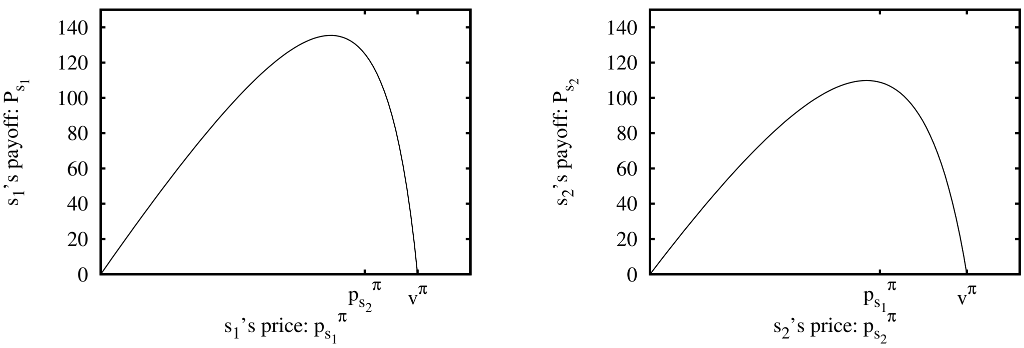

Figure 3.1: (*a*) Seller s_1's payoff function with one competitor and a population of ***spread buyers***, and (*b*) s_2's subsequent payoff function from s_1's best response.

Clearly, the best response price for s_1 is less than $p^\pi_{s_2}$; in fact in this instance it is 217.71. However, given this value as $p^\pi_{s_1}$ subsequently, s_2 is then faced with the payoff function illustrated in figure 3.1*b*. Of course, s_2 will respond to this value for $p^\pi_{s_1}$. Its best response is in this case 204.92. By using the sellers' payoff functions to iteratively calculate each seller's best response, this particular system is found to be at equilibrium when $p^\pi_{s_1} = p^\pi_{s_2} = 200$.

Clearly at this point the market share, and hence load, of each seller is also equal: $\vec{L_S} = \langle \frac{1}{2}, \frac{1}{2} \rangle$.

3.3 Conclusions and Discussion

In this chapter a model for computational resource allocation has been introduced, which uses the posted offer market mechanism to determine a mapping between resource providing and resource using nodes for the provision of the resource. A measurement of load on a resource providing node was also introduced. Notation was given to describe the seller and buyer populations and three buyer behaviour models were defined. Within this context, both buyer and seller payoff functions were described and the notions of buyer valuations and acceptable sellers were introduced.

The fundamental problem of achieving stable desired outcome resource allocations, under the assumption of self-interest and an absence of central or regional control was described. Although when these restrictions are not in place the problem is a trivial one, their presence makes the problem more challenging.

An initial analysis of the model with two competing sellers was conducted, considering populations of each buyer type. It was found that each of the three buyer types led the system to a different equilibrium, but that at each an evenly balanced load between the resource providing nodes was obtained.

The contributions of this chapter are therefore as follows:

- A model for computational resource allocation, using posted offer markets to determine to and from which nodes the resource is provided;
- A measurement of load on a resource providing node;
- Notation to describe seller and buyer populations, as well as the definition of three buyer behaviour models;
- Buyer and seller payoff functions, and notions of buyer valuation and the acceptability of sellers in the context of the model;
- An analysis of the equilibria of the system in the presence of three different buyer behaviour models, and a description of how this can lead to an evenly balanced load

as the outcome resource allocation.

The sellers' behaviour in this chapter can be considered to be myopically optimal, or *myoptimal* as described by Kephart et al. [9]. *Myoptimal* behaviour of this kind will result in the price dynamics described for each buyer behaviour being borne out as the system moves through time. Alternatively, agents endowed with foresight and capable of game theoretic reasoning would be expected to calculate the equilibrium positions identified and knowing that their payoff could not be bettered by any other action, adopt that price.

Of course, the assumptions required for either of these models of seller decision making are great. For *myoptimal* behaviour, at the very least accurate knowledge of the competitors' prices and buyers' strategies and private valuations would be required. For a fully game theoretic agent, knowledge of the other sellers' payoff functions and their strategies is also needed. For these reasons, as Kephart et al. [9] note, though this model of seller behaviour would be desirable from their perspective and also useful for the purposes of system-wide analysis, it is highly unlikely that it will be faithfully implemented, especially in larger systems. How then might sellers act instead in an attempt to emulate fully rational behaviour, and how might this lead a system to depart from the predicted outcomes described here?

CHAPTER 4

ACHIEVING A BALANCED LOAD WITH EVOLUTIONARY MARKET AGENTS

It is neither the strongest of the species that survive, nor the most intelligent,

but the one most responsive to change.

Charles Darwin

Chapter 3 introduced a model for decentralised computational resource allocation using self-interested agents engaged in a posted offer market mechanism. However, in the chapter's final section it was discussed that due to potentially unrealistic assumptions about the ability of sellers to obtain and process market information, the results of the analysis may misrepresent likely real world outcomes.

In this chapter sellers are modelled using an evolutionary approach, which requires them to possess neither global market information nor excessive computational processing power in order to adopt payoff maximising behaviour. Self-interested *evolutionary market agents* [1] are introduced, which act on behalf of resource providers to adaptively price their resources over time. The resource allocation model introduced in chapter 3 is then investigated in simulation when resource providing nodes use *evolutionary market agents* to price their resources. Simulation results for each of the three buyer behaviour types introduced in section 3.1.3 are presented and compared with their predicted results in section 3.2.

The contribution in this chapter is therefore to show that the use of online competitive co-evolution of self-interested resource providers can drive a decentralised market for those resources towards equilibrium, and to demonstrate that a balanced load emerges as a result. Furthermore, the system is robust to small changes in price. When buyers act as *spread buyers* this equilibrium is that predicted in the previous chapter's analysis and this is due to the smooth nature of these buyers' decision function. For *bargain hunters* and *time savers*, whose decision functions contain a step, *evolutionary market agents*' myopic nature leads the sellers' prices, and hence loads to oscillate rather than reach the Bertrand equilibrium, as would be expected of game theoretic agents. Mixed buyer populations are also investigated, with favourable results, and finally the scalability of the system is discussed and further simulation results are presented. These results were previously presented in [3].

Section 4.1 first highlights the reasons why the evolutionary simulation approach is necessary, in the light of the results from the previous chapter. Section 4.2 discusses potential approaches to developing an evolutionary algorithm and argues why that taken is appropriate in this scenario. Section 4.3 describes the *evolutionary market agent* algorithm used and section 4.4 updates some of the predictions from chapter 3 in light of this new approach. Section 4.5 describes the experimental set up and simulation results, including a comparison of the three buyer types and investigations into mixed populations of buyers and scalability. Section 4.6 discusses some design decisions taken in developing the evolutionary algorithm used, presenting comparative results with other potential variants, and finally section 4.7 concludes with a discussion.

4.1 Motivation

Chapter 3 introduced a model for decentralised computational resource allocation, using self-interested agents in a posted offer market mechanism. In this mechanism, selling agents advertise their resource at a price, and buying agents subsequently decide how much resource to purchase and from whom. As discussed in section 2.4, the key problem to be addressed by a seller in a posted offer market is that of how to adapt its price over time, such as to maximise its payoff. The game theoretic analysis in chapter 3 showed how when sellers behave either *myoptimally* or game theoretically, a stable equilibrium can be reached which results in an equal quantity of resource being provided by each seller.

However, a key problem facing sellers is that in only slightly larger systems than those explored in chapter 3, both *myoptimal* and game theoretic approaches becomes unfeasible, since they require both exact and global knowledge of the system's current state as well as a large amount of computational power [9]. Additionally, in the models investigated in this thesis, agents do not have access to the utility functions of other agents, as these are considered private information. As a result, they are not able to consider the payoff landscape with which they are faced, at least unless they employ some environment modelling technique, making use of information gained from interactions in the market. However, this is also a complex task, and troubled by the additional problem that with the effects of a highly dynamic system, the information learnt will be quickly out of date.

Kephart et al. [9] suggest that one possible alternative approach is to begin with some initial configuration of agents and use evolution to simulate their co-adaptation. However, the evolutionary method they propose still requires full knowledge of the payoff landscape, or else enough trial runs in the market to determine it, and therefore does not fully solve the scalability problem. This chapter takes a step further, making use of a lightweight evolutionary approach inspired by Price [111], in which agents co-evolve their market positions online, making use of only private information. This approach is taken both as a method to further analyse the behaviour of the model, especially with larger numbers of agents, and also as a prototype design for selling agents in a future embodiment of the model. A seller making use of this algorithm is termed an *evolutionary market agent.*

4.2 An Evolutionary Approach

With only private information available, sellers cannot directly calculate a best response price to the current state of the market, only sample the landscape using individual offers to see how they perform. One way for the seller to find the optimal market position, given the *black box* nature of the problem is therefore to take an exploratory approach. Evolutionary computation [68] provides an effective method of achieving this. Indeed, as section 2.2.4 discussed, evolution, or more accurately co-evolution when more than one agent is involved, has long been used to model competing players in repeated games. By using an evolutionary approach, sellers can expect to find good responses to the market. This approach allows us to achieve an approximation to the game theoretic outcome, under the assumption of private information.

Chapter 2 described a number of different approaches to using co-evolution in economic games, including the Michigan and Pittsburgh families of algorithms. Either approach could be used in developing an evolutionary algorithm to model adaptive pricing behaviour of sellers, and much literature exists comparing the suitability of each approach for other types of learning tasks [74].

A further consideration in the development of an evolutionary algorithm for use by an adaptive seller is that of strategy representation and the resulting search space. For example, key to the Pittsburgh family of learning algorithms is that each individual in the evolutionary algorithm's population contains a whole candidate rule set which defines the agent's behaviour. Individual rule sets should consider every scenario in which they agent may find itself, and prescribe an action accordingly. This is the approach taken by Cheung et al. [56] in their evolutionary modelling of petrol pricing in Australia.

Of course, the granularity of the rules' triggers can vary from the precise to the general, but as the size of the market increases, the number of rules is likely to also increase quickly. As Cheung et al. [56] point out, and as discussed in section 2.2.4, the complexity of the strategy search space may prove to be a major obstacle, due to the increased dimensionality and computational overhead required to find a good market position.

With this in mind, instead of encoding an explicit strategy capable of prescribing actions for all given market states, the approach taken here considers the problem of finding a good market position as a dynamic optimisation problem. Having each agent encode only their current offer, and by evolving these directly, the level of complexity is reduced considerably. The task before the evolutionary algorithm is then not to learn a *catch all* strategy, capable of determining the best action for any given future market state. Instead, it aims to find the optimal market position given the current state of the market and accept and deal with the notion that the market state may change during the optimisation process. In a sense, evolution itself is the strategy.

This approach has both positive and negative implications. Since the dimensionality of the problem is significantly reduced, the algorithm stands a good chance of finding a profitable market position within a very short space of time. The lower complexity also permits easier analysis of the fitness landscape associated with the search space, and hence provides further insight when selecting appropriate algorithms. Indeed, rather than being complex and high dimensional, the fitness landscape becomes the same as the agent's payoff landscape for the current market state.

However, relying on fitness information from the market itself, rather than from an offline model, means that each fitness evaluation is potentially very costly. Poor solutions must still be tested in the market, as are good solutions, in order to determine their fitness. In measuring the success of this approach, the key metric will not therefore be the fitness of the best individual, but the cumulative performance of the entire population over time.

4.3 The Evolutionary Market Agent Algorithm

An *evolutionary market agent* operates on behalf of a particular resource providing node in order to adaptively price its resource over time. Using evolutionary computation techniques, the agent evolves the market position of its host over time, in response to current market conditions. A population of prices is evolved online, with payoff information from the market being used as fitness values, and the objective being payoff maximisation. Competition between sellers is therefore driven by the co-evolution of their respective *evolutionary market agents*.

In this model a market position consists simply of price. Therefore each member of the population, an *individual* in evolutionary computation terminology, represents a real valued price. For each interaction in the market, an individual's price is adopted, and the resulting payoff provides its fitness.

The evolutionary algorithm used by an individual seller proceeds as in algorithm 1, and was first presented in [1].

Algorithm 1: The evolutionary market agent algorithm

1. Decide upon the design parameters to be used: initial price range $[p_{min}, p_{max}]$, population size and mutation factor, α. In the simulations described, $p_{min} = 0$, and $p_{max} = 500$. A population size of 20 was used, with a mutation factor, $\alpha = 0.1$.

2. Generate an initial population, Pop, and set $k = 1$. Each individual in Pop is a real value, drawn from the uniform random distribution $[p_{min}, p_{max}]$.

3. *Initial fitness testing*

   (a) Set the seller's offer to the value of the first individual in Pop, and enter the market for one market time-step. Record the seller's payoff, P_{s_i} as that individual's fitness.

   (b) Repeat for the next individual in Pop, until all initial individuals have been evaluated in the market.

4. *Probabilistic tournament selection*

   (a) Select four individuals, x_1, x_2, x_3 and x_4 from Pop, at random, such that $x_1 \neq x_2 \neq x_3 \neq x_4$.

   (b) Let champion c_1 be either x_1 or x_2, the fitness of whichever is greater with probability 0.9, the fitness of whichever is less otherwise.

   (c) Let champion c_2 be either x_3 or x_4, the fitness of whichever is greater with probability 0.9, the fitness of whichever is less otherwise.

5. Let the offspring, o, be a new individual with its price equal to the mid-point of c_1 and c_2.

6. Mutate o, by perturbing its value by a random number drawn from a normal distribution with mean zero and standard deviation α.

7. Insert o into Pop and remove either x_1 or x_2, whichever lost the first tournament, 4(b).

8. Set the seller's offer to the value encoded in o, and enter the market for one market time-step. Record the seller's payoff, P_{s_i} as o's fitness.

9. Repeat from step 4.

4.4 Predicted Outcomes with Evolutionary Market Agents

Most fundamentally to note about the *evolutionary market agent* algorithm is that, as expected, it has no knowledge of the buyers' decision functions, and hence does not calculate a best response to the current state of the market. The seller's behaviour is instead exploratory and myopic. This has important implications for the outcome of competition between *evolutionary market agents* in the presence of *bargain hunters* and *time savers*, since these buyer types have a step in their decision functions.

By contrast, *spread buyers* provide the benefit of a smooth decision function, and hence seller payoff function which degrades gracefully as a seller moves away from the optimal price. *Evolutionary market agents* lend themselves better to this behaviour. As an illustration, consider in the model described in chapter 3, the scenario when a seller s_1 is competing with s_2, where $v^\pi = 300$ and $p^\pi_{s_2} = 250$. Figure 4.1 illustrates the payoff s_1 can expect to receive as a function of its own next price, $p^\pi_{s_1}$, for the three buyer types we discussed in section 3.1.3.

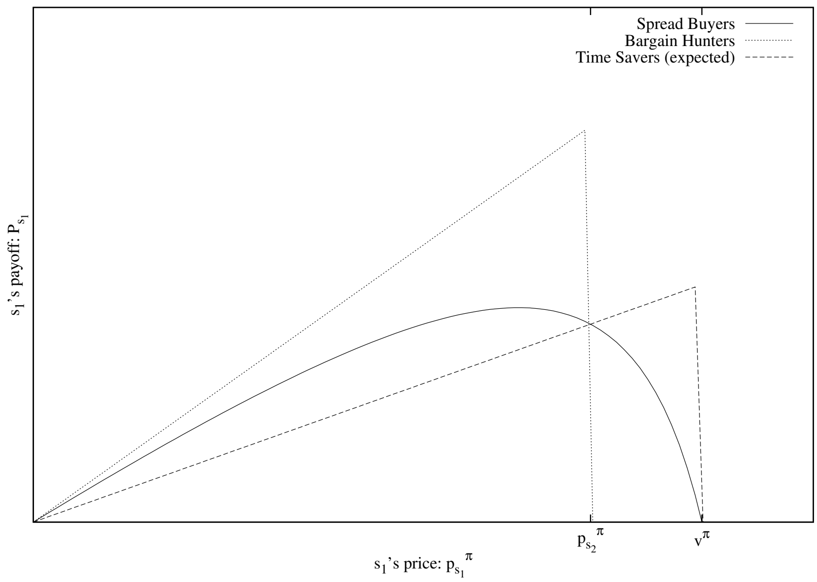

Figure 4.1: Seller s_1's payoff function with one competitor, for each of the buyer types. First presented in [3].

The step function is clearly visible here, for sellers dealing with either *bargain hunters* or *time savers*. In the absence of full knowledge of the buyers' decision functions (i.e. the curves shown here), and relying upon a myopic exploratory heuristic, a seller would be expected to oscillate between receiving a near-optimal payoff, and one of zero.

Such an oscillation between high and low payoffs is not conducive of stable outcomes. Indeed, for *bargain hunters*, the position of the step for a given seller will depend upon its competitors' prices. Since the same is also true of the other competitors, the exact equilibrium position of the step, given by the sellers' prices, will be fairly arbitrary.

For *spread buyers* however, a small move away from the optimal price leads to a small change in payoff for the seller. The resulting impact upon the other sellers' payoff functions is also minimal. This incremental characteristic of a market of *spread buyers* leads to a less brittle, more predictable system.

4.5 Simulation Results

Since the co-evolutionary approach aims to emulate rational *myoptimal* behaviour under the assumption of only private information, the hope is that the system reaches an approximation to the predicted outcome in chapter 3. However, given the *evolutionary market agents*' stochastic, exploratory nature, statistical results gained experimentally are useful. The scenarios described in section 3.2, and some larger, more complex examples are now explored in simulation.

4.5.1 A Baseline Scenario

Figure 4.2 illustrates the evolution of this behaviour over time in a scenario with two resource providing nodes, such that $S = \{s_1, s_2\}$. Both s_1 and s_2 are represented by independent *evolutionary market agents*, as described in section 4.3, each with a population size of 20 and a mutation factor of 0.1. B consists of ten homogeneous *spread buyers*, such that $v^{\pi} = 300$.

Figure 4.2 clearly shows the ability of the approach to achieve a roughly even load between the two resource providers in a short time. This is due to the evolutionary agents' competitively co-evolving their prices to within close proximity of each other quickly, result-

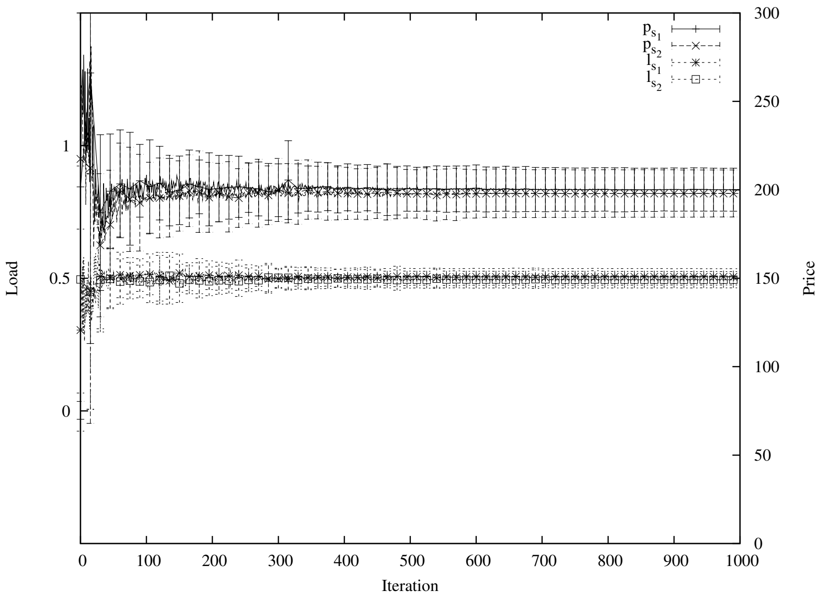

Figure 4.2: Evolution of price (above) and load (below) of two resource providing nodes over time, with a population of 10 *spread buyers*. A stable evenly balanced load emerges. Mean and standard deviation over 30 independent runs. First presented in [3].

ing in roughly even shares of the market. Following these exploratory fluctuations, the loads stabilise as the prices converge to the predicted equilibrium, here 200. At this point, the loads are highly equal. Due to diverse populations within each agent's population however, their prices, and hence also the allocation of resources continue to vary to a small degree.

4.5.2 Comparing Buyer Types

Let us now turn our attention to the behaviour of the system in the presence of *bargain hunters* and *time savers*. Intuitively, from section 4.4 above, oscillatory behaviour would be expected when these buyers' decision functions, with their step characteristics, are present. When observing the load on resource providing nodes in simulation this unreliable behaviour, along with a high standard deviation, is indeed manifested.

Figures 4.3 and 4.4 illustrate this with two sellers and a population of 10 *bargain hunters*, and *time savers* respectively. The behaviour illustrated above in figure 4.2, for *spread buyers*, is clearly more desirable.

The presence of the step in the decision function for both *bargain hunters* and *time savers*

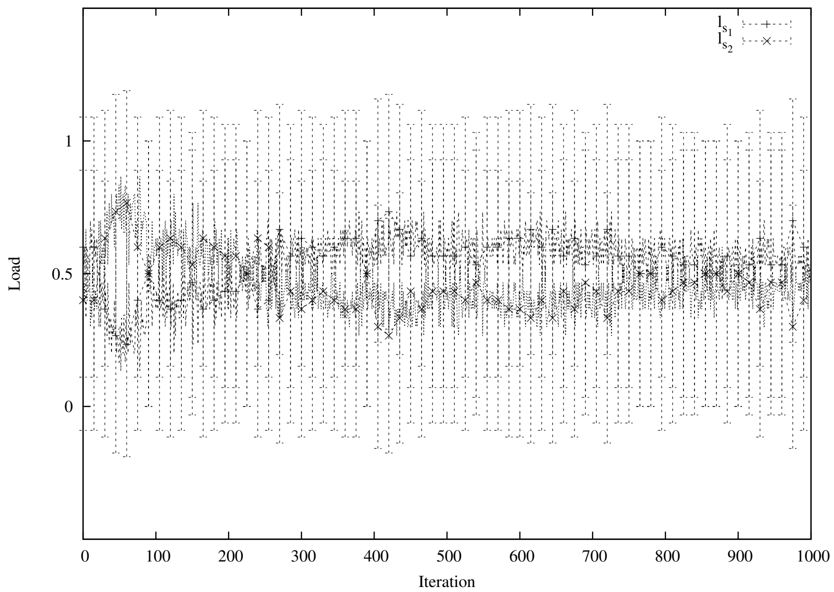

Figure 4.3: Evolution of the load on two resource providing nodes over time, with a population of 10 *bargain hunters*. The step in the buyers' decision function leads to oscillations and unpredictability. Mean and standard deviation over 30 independent runs. First presented in [3].

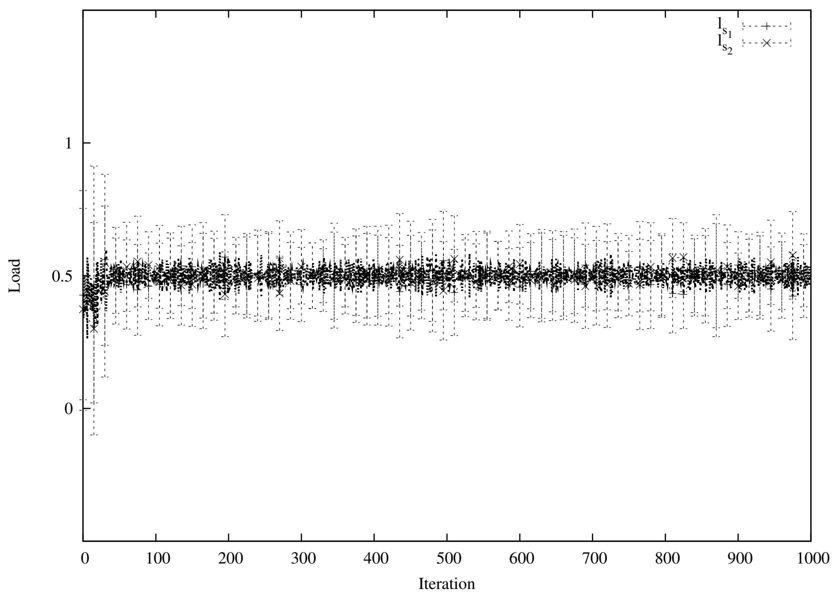

Figure 4.4: Evolution of the load on two resource providing nodes over time, with a population of 10 *time savers*. The step in the buyers' decision function has less of an oscillatory effect. Mean and standard deviation over 30 independent runs. First presented in [3].

clearly leads to oscillations in the load allocated between the resource providing nodes; the allocations do not remain stable over time. The graceful degradation of the *spread buyers*' decision function, however, leads to a highly stable and evenly balanced load.

4.5.3 Mixed Buyer Populations

As has been shown, one benefit of *spread buyers*' behaviour is the graceful degradation of a seller's market share as its price moves away from the optimum (or the optimum moves away as competitors update their prices). This smooth curve allows the sellers' evolutionary algorithms to easily find the new optimum price. By contrast, as we saw in section 4.5.2, buyer behaviours with a step in their decision function lead to more erratic, less stable equilibria.

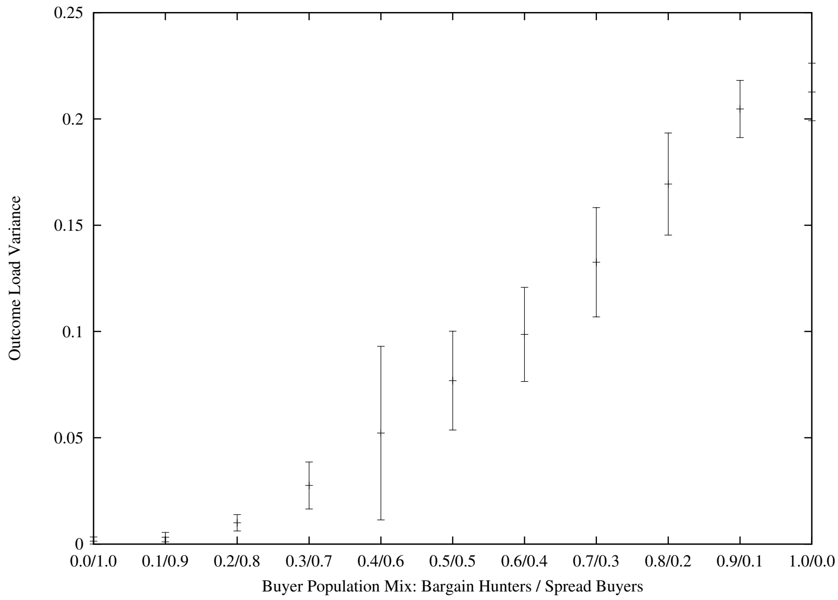

Figure 4.5: The proximity of the outcome resource allocation to an evenly balanced load degrades gracefully with respect to the proportion of *bargain hunters* in a population otherwise consisting of *spread buyers*. Each data point is from 30 independent runs of a simulation with two sellers and 100 buyers, which vary along the x axis. First presented in [3].

However, how disruptive is the presence of such buyer behaviour in a population otherwise consisting of *spread buyers*? The experimental result in figure 4.5 indicates a graceful degradation of performance in the presence of an increasing proportion of *bargain hunters* in the population.

4.5.4 Scalability

In addition to the unrealistic assumptions outlined in section 4.1, one of the key reasons for adopting this online evolutionary approach rather than attempting to implement *myoptimal* sellers or a full strategy learning algorithm is scalability. Since information about buyers and competitors does not need to be obtained, and the payoff landscape does not need to be computed by the seller, the complexity of pricing algorithm itself is not dependent on market size. Additionally, due to the distributed, decentralised nature of the mechanism employed, the approach scales well.

This scalability leads to the behaviour demonstrated in the above simple scenario also being observed in much larger resource markets. Figure 4.6 illustrates the evolution of price and load for $1,000$ sellers and $10,000$ buyers in a typical run of the simulation. Figure 4.7 shows the mean and standard deviation of the load variance for 30 independent runs. It is clear that in this respect the approach is highly scalable, as results are of a similar form to the smaller simulation. Figure 4.8 shows the time taken for the system to reach within 1% of the predicted outcome allocation, from the initial allocation, as the number of sellers increases.

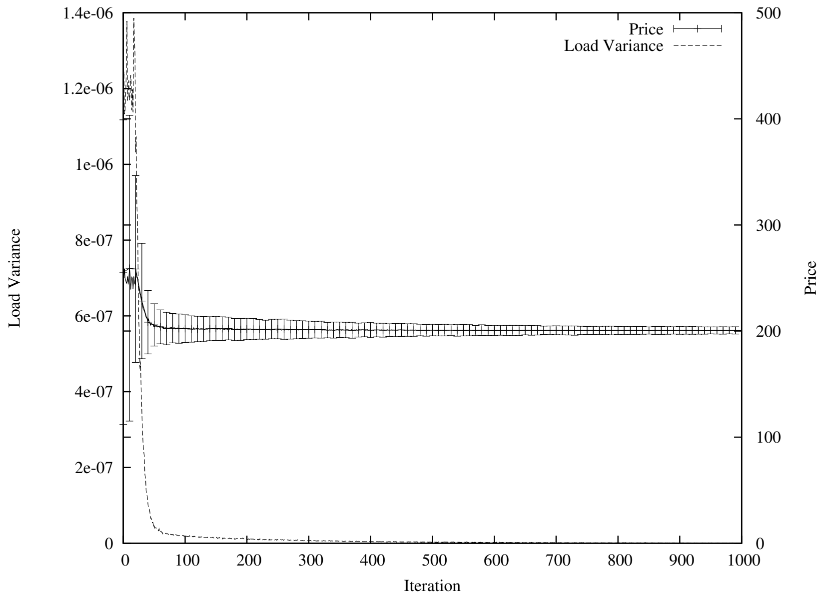

Figure 4.6: Evolution of price and load variance between $1,000$ resource providing nodes over time, with a population of $10,000$ *spread buyers*, demonstrating a high degree of scalability. Results from a typical run. First presented in [3].

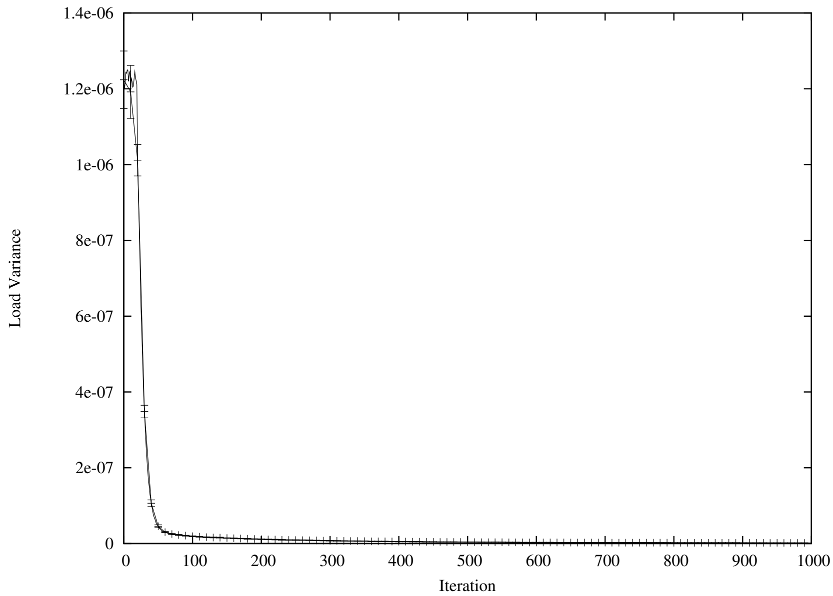

Figure 4.7: Load variance between $1,000$ resource providing nodes over time, with a population of $10,000$ *spread buyers*, demonstrating a high degree of scalability. Mean and standard deviation over 30 independent runs. First presented in [3].

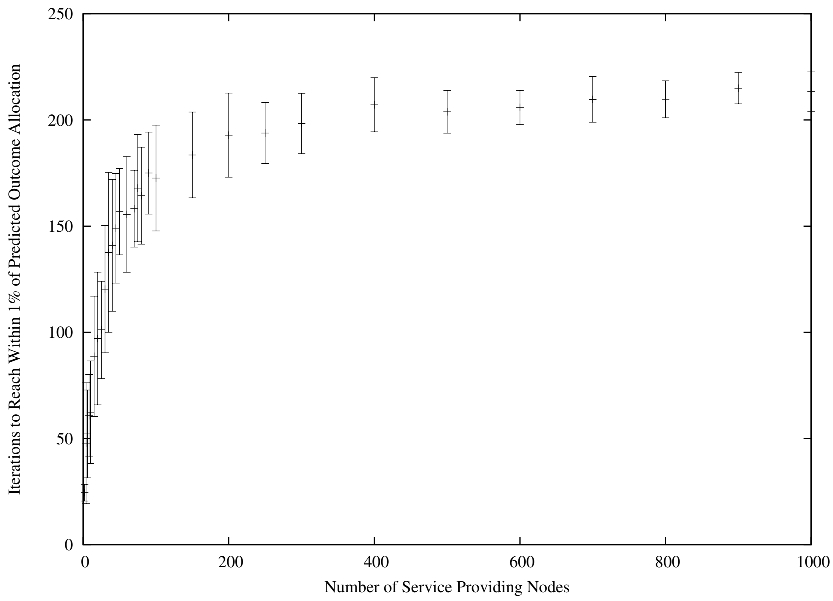

Figure 4.8: Time taken to reach close proximity to the predicted outcome allocation (within 1%). Each data point represents the mean and standard deviation over 30 independent runs. First presented in [3].

A further important question concerning the scalability of the approach, and one that is shared with many other market-based mechanisms such as CATNET [22], is that of how knowledge of offers might be disseminated around the system. In some embodiments of the model, such as networks making use of wireless technology, broadcasting may indeed be achievable. In peer-to-peer networks, the broadcast facility may be emulated through algorithms such as flooding [113], distributed hash tables [114] or epidemic algorithms [115]. However, Ardaiz et al. [22] argue quite correctly that decentralised and self-interested resource allocation approaches should prefer decentralised, incentive-compatible discovery mechanisms.

4.6 Design of the Evolutionary Algorithm

The evolutionary algorithm used here is not the result of extended experimental tuning or design, and as such it is highly likely that it may be improved upon. This thesis does not claim that the *evolutionary market agents* presented here is superior to other already existing adaptive pricing algorithms, however this is not the primary concern at this stage. Rather, this chapter emphasises the sufficiency of the online co-evolutionary approach to drive competition between sellers and achieve the effects described. A useful extension to this work would be to perform a comparison with other suitable strategies, such as the Gjerstad-Dickhaut [116] and Roth-Erev [117] algorithms.

However, during the course of the experimentation, some results concerning the evolutionary algorithm itself were obtained which are now reported. Particularly, the decision to use either stochastic or deterministic selection and replacement operators affects the algorithm's performance. Two variants of each operator were tested. For selection, the variants were the stochastic selection method described in section 4.3 and a more deterministic version where the winner of each tournament was always chosen to be a parent. For replacement, the variants were the deterministic operator described in section 4.3, where the loser of the first tournament was always replaced, and an alternative where with a probability of 0.1 the individual replaced was instead selected at random from all those in the two tournaments. Combining these, the four variants of the algorithm tested were as follows:

- Both deterministic selection and replacement operators,

- Deterministic selection but probabilistic replacement,
- Probabilistic selection and probabilistic replacement, and
- Probabilistic selection but deterministic replacement.

In the experiments, all variants of the algorithm were able to find the equilibrium price, but differed on their behaviour once at the optimum. Figure 4.9 shows the standard deviation of each algorithm's price about the mean for a seller over 30 independent runs.

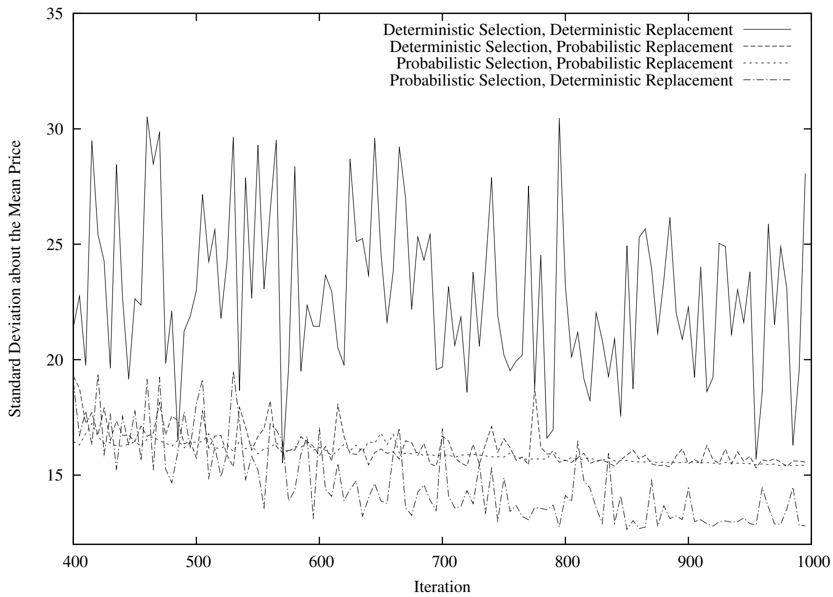

Figure 4.9: Standard deviation about the mean price for four algorithm variants. Probabilistic selection combined with the deterministic replacement operator enables the algorithm to remain closer to the equilibrium price. First presented in [3].

As can be seen, the use of both deterministic operators led to a high and highly volatile standard deviation, indicating frequent significant fluctuations in price. The two variants with probabilistic replacement operators performed very similarly to each other, with lower and significantly more stable standard deviations. The lowest standard deviation was achieved by the variant with the probabilistic selection but deterministic replacement operators. This is the algorithm described in section 4.3 and it is for this reason that it was chosen for the majority of the simulation experiments in the thesis.

Intuitively, probabilistic operators could be expected to improve performance. This is since, in the early part of the simulation, a price higher than the equilibrium might yield a high payoff for a seller, but once the other competitors had responded to this, it would no longer be possible for any seller to reach such a high payoff. This price would remain in the population, and frequently win tournaments, despite its adoption at a later iteration being unlikely to reproduce the high payoff. The market would have moved on and its fitness would be out of date. Probabilistic selection ensures that other individuals, with lower fitness values but perhaps more suited to the current market conditions, can get selected instead. It is interesting then that probabilistic replacement did not bring about a similar performance gain. However, though it might be useful to occasionally select individuals with only an average fitness in future iterations, it is likely that individuals with very low fitness would perform poorly in any scenario, and there would therefore be little to gain by retaining them. A more formal analysis of the populations over time will prove illuminating.

4.7 Conclusions and Discussion

This chapter recognises that while they are useful to analyse the equilibria in the model, the *myoptimal* and game theoretic approaches taken by sellers in chapter 3 rely on a number of unrealistic assumptions. Firstly, that in order for a seller to identify its immediate best response offer to a given market scenario, it must first have access to every competitor's current offer as well as knowledge of the buyers' decision functions. Secondly, it assumes that each seller has sufficient computational power to calculate the resulting optimum. For a fully game theoretic seller, the additional knowledge of its competitors' pricing algorithms is also required.

Given these unrealistic assumptions, this chapter attempts to provide an alternative approach, where sellers each use an evolutionary algorithm to explore the payoff landscape and optimise their market position. It was shown that this co-evolution of offers can lead the system to approximate the outcome resource allocations predicted in chapter 3. Indeed, co-evolution itself is the competitive force which drives the market towards equilibrium.

The relationship between experimental performance in simulation, and game theoretic predictions for the three buyer behaviour models defined in section 3.1.3 was explored. Buyer behaviours with a step in their decision function were found to lead to unstable, erratic outcome allocations when sellers' prices were set by *evolutionary market agents*. Those buyers with smooth decision functions however, such as *spread buyers*, lead to a highly even and stable balanced load, which is also being robust to small changes in price. It was further shown that this performance degrades gracefully as the proportion of buyers with step functions in the population increases.

Therefore, the contributions of this chapter are as follows:

- An argument that *myoptimal* and game theoretic selling agents are unlikely to be faithfully implemented in decentralised computational systems where nodes are self-interested, since they require information about other agents which is typically considered private and unavailable;
- The description of an alternative approach, using *evolutionary market agents*, which sellers may use to attempt to optimise their prices within the constraints of the information available.
- Experimental evidence from simulation runs that when sellers use *evolutionary market agents*, the system is able to approximate the game theoretically predicted equilibrium, and hence evenly balance the load across the resource providing nodes, when buyers adopt the *spread buyer* behaviour.
- Experimental evidence that due to a step in the decision functions of *bargain hunter* and *time saver* buyers, the system does not converge to the predicted equilibrium and instead remains unpredictable and unstable in these cases;
- Further experimental evidence showing that when buyer populations are mixed, the system's performance degrades gracefully as the proportion of buyers with steps in their decision functions increases, in a population otherwise consisting of *spread buyers*.

- Further experimental evidence which demonstrates that the system scales well with respect to the number of agents participating in the system and a balanced load can be achieved with very large numbers of agents.
- Analysis of experimental results for four different variants of the *evolutionary market agent* algorithm, with the conclusion that probabilistic selection and deterministic replacement operators provide good performance.

As discussed in section 3.2, the underlying theoretical idea which describes how the system achieves a balanced load is that of Bertrand competition. However, it should be noted Bertrand competition relies on the presence of a number of additional potentially unrealistic assumptions [34]. Three of these are of particular interest. Firstly, Bertrand competition assumes no collusion between sellers. Secondly, it is also assumed that sellers only compete on price, and are otherwise unable to differentiate their products in the market. Thirdly, sellers are assumed to have the same marginal cost.

Here, collusion is ruled out by the model itself, but it would be remiss to ignore the possibility of its occurrence in real world embodiments. Cheung et al [56] showed in a similar evolutionary simulation that if sellers are able to reliably predict the responses of their competitors, then even without knowledge of their pricing algorithms, they may implicitly collude in order to raise prices and hence their payoffs. This behaviour is not observed in the simulations here, since the sellers do not retain historical information concerning each other. It is also unlikely, as noted by Cheung et al. that this collusion would occur in very large scenarios with many competing sellers. However, in smaller oligopolistic markets involving strategic sellers, the possibility of this affecting the outcomes should be borne in mind.

It is also likely in real world embodiments that truly self-interested sellers would attempt to violate the second assumption of Bertrand competition listed above, by differentiating their product other than on price alone. Particularly in the presence of heterogeneous buyers, factors such as quality of service or branding could provide targeted product differentiation, which can lead to increased profits for the seller and alternative overall outcomes [34]. Looking ahead, chapters 6 and 7 therefore investigate the effect of considering additional quality attributes over which to describe the resource, such that sellers are able to achieve such differentiation.

Finally, the balanced load outcome presented here is heavily reliant on the third assumption, that of equivalent marginal costs between sellers. This is intentional, since the approach is motivated by large homogeneous resource providing environments, where it does not matter from whom the resource is provided. It is of course reasonable to expect that should the resource providing nodes not be homogeneous, and instead vary in their ability to provide the resource which is being traded, then a different outcome resource allocation would be obtained. This important topic is the subject of chapter 5.

CHAPTER 5

ACHIEVING ADDITIONAL OBJECTIVES IN RESOURCE ALLOCATION

I may not have gone where I intended to go, but I think I have ended up where I needed to be.

Douglas Adams

This chapter builds upon the previously studied market-based approach to achieving resource allocation in decentralised systems, by considering heterogeneous provider nodes. In this scenario, providers may be said to value their resources differently. By considering sellers' heterogeneous costs and valuations, an extended version of the model described in chapter 3 is able to bring about additional stable uneven outcome resource allocations.

Here it is demonstrated how, given such valuations and a model of buyer behaviour, the outcome allocation may be predicted. Furthermore, a method is outlined for determining sellers' parameters in order to use the approach to achieve a desired stable objective. The behaviour of the extended model is analysed both game theoretically and in simulation using *evolutionary market agents*. As with chapters 3 and 4, the impact of different buyer behaviour models is considered. Importantly, particular desired resource allocations are achieved transparently to resource users, as no modification to the buyers is required.

Section 5.1 begins the chapter by motivating the need to extend the original model. Section 5.2 then describes the extension, paying particular attention to changes needed to the sellers' payoff functions in order to encapsulate their heterogeneity. Section 5.3 analyses the newly extended model game theoretically, considering the impact of buyers behaving as *bargain hunters*, *time savers* and *spread buyers*. The problem of determining parameters for sellers in order to achieve desired outcome resource allocations is addressed in section 5.4, before experimental results in section 5.5 demonstrate the achievement of the predicted outcomes when *spread buyers* and *evolutionary market agents* are used in simulation. Section 5.6 characterises the impact of varying the sellers' valuations with respect to each other, and illustrates why issues of stability may arise in monopolistic scenarios. Finally, section 5.7 concludes with a discussion.

5.1 Motivation

So far, this thesis has proposed an evolutionary market-based approach to achieving an evenly balanced load between resource providing nodes in a fully decentralised system where agents do not cooperate. The approach, which makes use of an artificially created posted offer market for the resource, relies purely upon self-interest, and no individual node has any desire in favour of the balanced outcome. In the mechanism, resource providing nodes quote prices for equivalent resources and resource users select how much and from whom to purchase. *Evolutionary market agents*, acting on behalf of resource providing nodes, evolve the price quoted by their respective seller over time. The payoff from the live market is used as the fitness function for each seller's evolutionary algorithm.

However, the approach described in chapters 3 and 4 considers that all resource providing nodes have an equivalent ability to provide the resource. In many real world scenarios this will not be the case. For example, one node may be able to perform more work than another in the same amount of time.

In this kind of scenario, with heterogeneous resource provider abilities, the approach described so far is not able to take account of these factors. Its use will result in an even resource allocation between the resource providing nodes, regardless of their heterogeneous abilities. This is therefore likely to be an inefficient outcome. In this chapter the approach is extended in order to take advantage of this information in providing an appropriate, stable and desired outcome allocation.

5.2 Heterogeneous Resource Providing Nodes

As in the homogeneous model introduced in chapter 3, consider a set of resource providing nodes, S, each member of which provides a functionally equivalent, quantitatively divisible resource π, which may vary only in price. Imagine then a large population of service users or buyers, B, each member of which aims to consume one unit of the resource π, at regular intervals.

In chapters 3 and 4, the objective was to balance the load evenly, such that all the resource providers in S provided an equal amount of π across the population of resource users. In this chapter however, the objective is to predict what outcome resource allocations will be reached, given heterogeneous seller abilities. These abilities are reflected in the model through the sellers' own private valuations for the resource π. Furthermore, this chapter attempts to understand how the approach may be extended in order to achieve a particular chosen stable outcome resource allocation. As with the homogeneous case, while this is indeed a trivial problem when cooperation between agents may be assumed, the objective here is to achieve this using self-interest, in a fully decentralised system. Additionally, the possibility of complex buyer decision functions means that there may not be a straightforward mapping between valuations and the outcome allocation.

The market mechanism and notation used are the same as described in chapter 3. In particular, recall that a particular configuration for the provision of the resource in the model, such as the outcome resource allocation may be expressed by the vector $\vec{L_S} = \langle l_{s_1}, l_{s_2}, ..., l_{s_n} \rangle$, where $n = |S|$. As described in section 3.1, for ease of ease of comparison between scenarios, the values in the vector are often presented normalised by the total quantity of resource being provided.

In this way, any desired outcome allocation may be described. As we have seen, an evenly balanced load for example may be written as $\langle \frac{1}{n}, \frac{1}{n}, ..., \frac{1}{n} \rangle$. Similarly, an allocation between six providing nodes, where the even numbered nodes provide exactly twice the resource as the odd numbered nodes would be $\langle \frac{1}{9}, \frac{2}{9}, \frac{1}{9}, \frac{2}{9}, \frac{1}{9}, \frac{2}{9} \rangle$.

In the original version of the model all resource providing nodes were implicitly assumed to have an equal ability to provide the service. The objective was to balance the usage or load equally across them. In certain realistic embodiments however, it is unlikely that this will be the case. Leaving aside any consideration of differences in the overall *capacity* of resource providing nodes, that is to say the total quantity of π which a single node may provide, here a heterogeneity of the *ability* of nodes in the population to provide the resource is considered. One resource providing node, for example, may be able to perform the same task in a shorter amount of time than another, due to an increased capability. The original payoff function takes no account of this.

Such heterogeneities of resource providing nodes' abilities may be represented by their selling agents using a notional cost of provision, along with different relative valuations of π when compared with this cost. Since, functionally speaking, equivalent quantities of work are substitutable irrespective of which node performed them, it might be expected that a node able to provide more work than another for the same cost would place a lower private valuation on performing that quantity of work, when compared with the cost.

This valuation may be built in to the existing model in the form of a sellers' payoff function which takes account of the cost of the provision of π, $c^{\pi}_{s_i}$, and a preference weight, w_{s_i}, on the price.

In the original model described in section 3.1, seller s_i's utility gain, or payoff, P_{s_i}, was given by

$$P_{s_i} = \sum_{j=1}^{|B|} p^{\pi}_{s_i} q_{ij} \, , \tag{5.1}$$

or indeed

$$P_{s_i} = p^{\pi}_{s_i} \times l_{s_i} \, . \tag{5.2}$$

Introducing a node's cost of providing π, $c^{\pi}_{s_i}$, and taking account of a utility preference weight, w_{s_i}, s_i's payoff can instead be proposed to be as follows:

$$P_{s_i} = \sum_{j=1}^{|B|} q_{ij} (w_{s_i} p^{\pi}_{s_i} - c^{\pi}_{s_i}) \, , \tag{5.3}$$

or

$$P_{s_i} = (w_{s_i} p^{\pi}_{s_i} - c^{\pi}_{s_i}) \times l_{s_i} \, . \tag{5.4}$$

This extended version of the model was first presented and analysed in [2].

5.3 Predicted Outcomes

As was done in the homogeneous case in chapter 3, by considering agents' best responses in attempting to maximise their instantaneous payoff, it is possible to identify potential equilibria and hence predicted outcome resource allocations.

5.3.1 A Baseline Example

Let us begin by exploring the effect of heterogeneous seller abilities by means of a two node example in the presence of a population of *spread buyers.* Recall that for such buyers, the quantity of π purchased by buyer b_j from seller s_i,

$$q_{ij} = \frac{(v_{b_j}^{\pi} - p_{s_1}^{\pi})}{(v_{b_j}^{\pi} - p_{s_1}^{\pi}) + (v_{b_j}^{\pi} - p_{s_2}^{\pi})}, \tag{5.5}$$

where $v_{b_j}^{\pi}$ is the valuation, and hence maximum price payable by each buyer.

When faced with a population of *spread buyers*, the payoff functions (from equation 5.3) for s_1 and s_2 are therefore

$$P_{s_1} = \sum_{j=1}^{|B|} \frac{v_{b_j}^{\pi} - p_{s_1}^{\pi}}{(v_{b_j}^{\pi} - p_{s_1}^{\pi}) + (v_{b_j}^{\pi} - p_{s_2}^{\pi})} (w_{s_1} p_{s_1}^{\pi} - c_{s_1}^{\pi}), \tag{5.6}$$

and

$$P_{s_2} = \sum_{j=1}^{|B|} \frac{v_{b_j}^{\pi} - p_{s_2}^{\pi}}{(v_{b_j}^{\pi} - p_{s_2}^{\pi}) + (v_{b_j}^{\pi} - p_{s_1}^{\pi})} (w_{s_2} p_{s_2}^{\pi} - c_{s_2}^{\pi}), \tag{5.7}$$

Sellers s_1 and s_2 will then attempt to maximise their respective payoff function as before. The outcome resource allocation is taken to be that which occurs when the system is at equilibrium.

In this example, if s_1 represents a node with a greater ability to provide the resource than s_2, then we might say that $w_{s_1} = 1.0$ and $w_{s_2} = 0.5$; s_2 values its work twice as much when compared to money. For the purposes of the example it is initially assumed that both sellers bear a fixed cost; $c_{s_1}^{\pi} = c_{s_2}^{\pi} = 100$, and that $v_{b_j}^{\pi} = 300$, to give a reasonable range of prices.

Using the iterative procedure described in section 3.2.3, it can be calculated that the system in this example is at equilibrium when $p^{\pi}_{s_1} = 243.1$ and $p^{\pi}_{s_2} = 262.4$. At this point, neither s_1 nor s_2 may increase its respective payoff by unilaterally choosing a different price.

The allocation at these prices therefore, is the expected outcome allocation. This may be calculated for each resource provider by substituting equation 5.5 into equation 3.1:

$$l_{s_i} = \sum_{j=1}^{|B|} \frac{(v^{\pi}_{b_j} - p^{\pi}_{s_1})}{(v^{\pi}_{b_j} - p^{\pi}_{s_1}) + (v^{\pi}_{b_j} - p^{\pi}_{s_2})} \, . \tag{5.8}$$

Solving this using the sellers' weights, the determined equilibrium price values and assumed costs in the example,

$$l_{s_1} = 0.6 \times |B| \, , \tag{5.9}$$

and

$$l_{s_2} = 0.4 \times |B| \, . \tag{5.10}$$

In other words, we can expect s_1 to take 60% of the load, and s_2 to take 40%, given the above valuations and costs.

This method for determining the outcome resource allocation may easily be generalised to other buyer behaviours, by replacing equation 5.5 with one which captures the appropriate buyer behaviour.

5.3.2 Some More Complex Examples

Now let us consider a more complex case for each of the three buyer behaviours investigated in chapter 3.

Bargain Hunters

Recalling results from chapter 3, game theoretic and *myoptimal* sellers in the presence of *bargain hunters* are expected to compete downwards on price, until they reach their reserve valuation for the resource. Since in the homogeneous case the sellers' costs were the same, this meant that their prices at equilibrium were also the same.

In the heterogeneous case, though competition on price will still occur, critically the lower limit price for each seller will depend upon its cost and preference weight. As might be expected, that seller (or those sellers) which are able to offer the lowest price will take the entire market, leaving those unable to compete with nothing.

This may be illustrated this by means of a three node example, such that $S = \{s_1, s_2, s_3\}$, where $c^{\pi}_{s_1} = c^{\pi}_{s_2} = 100$, $c^{\pi}_{s_3} = 150$ and $w_{s_1} = w_{s_2} = w_{s_3} = 1$. We therefore have the following payoff functions:

$$P_{s_1} = (p^{\pi}_{s_1} - 100) \times l_{s_1} \,, \tag{5.11}$$

$$P_{s_2} = (p^{\pi}_{s_2} - 100) \times l_{s_2} \,, \tag{5.12}$$

and

$$P_{s_3} = (p^{\pi}_{s_3} - 150) \times l_{s_3} \,. \tag{5.13}$$

Considering first s_1, clearly if its offer is to be accepted by any buyer, indicated by a positive load, $l_{s_1} > 0$, a non-negative payoff is obtained if and only if $p^{\pi}_{s_1} \geq 100$. The same is true respectively of s_2. However, for s_3 this is the case if and only if $p_{s_3} \geq 150$. s_3 is therefore unable to compete with s_1 and s_2 and will discontinue participation in the price war once the best price drops below 150. s_1 and s_2 however, will continue undercutting each other until their prices reach 100, at which point any further price cut would leave them with a negative payoff themselves. However, since they both share the same limit price, they will remain at equilibrium with $p_{s_1} = p_{s_2} = 100$, and therefore $\vec{L_S} = \langle \frac{1}{2}, \frac{1}{2}, 0 \rangle$.

Now let us add a fourth node to S, s_4, such that $c^{\pi}_{s_4} = 150$ and $w_{s_4} = 2.5$. s_4 has the following payoff function:

$$P_{s_4} = (2.5 \times p^{\pi}_{s_4} - 150) \times l_{s_4} \,. \tag{5.14}$$

Now, s_4's cost is also 150, as with s_3, however since its preference weight is 2.5, its payoff will not drop below zero for positive loads so long as $p_{s_4} \geq 60$. This is clearly a more competitive position than s_1 and s_2 are able to take. The expected allocation at equilibrium for this four node system is therefore $\vec{L_S} = \langle 0, 0, 0, 1 \rangle$.

More generally, for purchases made by hyperrational *bargain hunters*, the seller or sellers able to offer the lowest price will share the load evenly between them. Any unable to offer this price will have a load of zero. The limit price for a seller s_i is $\frac{c^{\pi}_{s_i}}{w_{s_i}}$.

This example enables us to observe the *all or nothing* nature of a population of buyers made up entirely of *bargain hunters*, and it is here that the limit of Bertrand competition in its classic sense is reached. In order to elicit more complex resource allocations, we must turn to other buyer behaviours.

Time Savers

A buyer population consisting wholly of *time savers* will lead to a different outcome. Recall that in the homogeneous case sellers' best response prices were dependent not on each other, but on the buyers' reserve price. So long as a seller ensures that it is in the buyers' set of acceptable sellers, S_B, it can maximise its payoff simply by maximising its price within that constraint. The sellers' cost and valuation parameters do not come into play in this case, other than to ensure that there exists an offer which places it in S_B which does not return a negative payoff for the seller itself.

Therefore, the result for the homogeneous case for *time savers* also holds where sellers have heterogeneous costs and preference weights. Since in this case each seller's market share is independent of its price, given a seller's membership of S_B, any operation on the price will not affect the load at equilibrium.

Spread Buyers

As was demonstrated in section 5.3.1, the iterative method used in chapter 3 for *spread buyers* may be also employed in more complex cases to determine equilibria, and hence outcome resource allocations. For the three node example discussed in section 5.3.2, the sellers' payoff functions are now

$$P_{s_1} = \sum_{j=1}^{|B|} \frac{v^\pi - p^\pi_{s_1}}{3v^\pi - (p^\pi_{s_1} + p^\pi_{s_2} + p^\pi_{s_3})} (w_{s_1} p^\pi_{s_1} - c^\pi_{s_1}) \,, \tag{5.15}$$

$$P_{s_2} = \sum_{j=1}^{|B|} \frac{v^\pi - p^\pi_{s_2}}{3v^\pi - (p^\pi_{s_1} + p^\pi_{s_2} + p^\pi_{s_3})} (w_{s_2} p^\pi_{s_2} - c^\pi_{s_2}) \,, \tag{5.16}$$

and

$$P_{s_3} = \sum_{j=1}^{|B|} \frac{v^\pi - p^\pi_{s_3}}{3v^\pi - (p^\pi_{s_1} + p^\pi_{s_2} + p^\pi_{s_3})} (w_{s_3} p^\pi_{s_3} - c^\pi_{s_3}) \,. \tag{5.17}$$

Following the iterative method, it can be calculated that this system is at equilibrium when $p^\pi_{s_1} = p^\pi_{s_2} = 221.75$ and $p^\pi_{s_3} = 237.49$. At this point, none of the sellers may increase their respective payoff by unilaterally choosing a different price. The allocation at these prices is our predicted outcome allocation. This may be calculated for each resource provider using equation 5.4:

$$l_{s_i} = \sum_{j=1}^{|B|} \frac{v^\pi - p^\pi_{s_i}}{3v^\pi - (p^\pi_{s_1} + p^\pi_{s_2} + p^\pi_{s_3})} \,. \tag{5.18}$$

In the example discussed here, we therefore have an outcome resource allocation of $\vec{L_S} = \langle 0.3575, 0.3575, 0.2850 \rangle$. In other words, given the above valuation and cost parameters, we expect s_1 and s_2 to take 35.75% of the load each, and s_3 to take 28.5%.

Adding in the fourth seller at this point, the calculations must be repeated. In the four node example discussed in section 5.3.2, but with *spread buyers*, the prices at equilibrium are $p^\pi_{s_1} = p^\pi_{s_2} = 214.53$, $p^\pi_{s_3} = 233.25$ and $p^\pi_{s_4} = 200.73$. The outcome resource allocation is therefore $\vec{L_S} = \langle 0.2536, 0.2536, 0.1981, 0.2946 \rangle$.

Critically, this result demonstrates that while *bargain hunters* and *time savers* do not provide the necessary tools with which to achieve stable uneven resource allocations, *spread buyers* by contrast do just that. Given self-interested sellers with the appropriate cost and value parameters, a population of *spread buyers* will make self-interested decisions such that predictable outcome allocations are obtained. This is without any modification to the buyers' behaviour, and indeed without any desire or additional knowledge required on their part. These results were presented in [2] and [3].

5.4 Determining Cost and Valuation Parameters

The ability to predict an outcome allocation, the allocation at equilibrium, given an initial configuration of nodes is in itself useful. However, this thesis is concerned with resource allocation tasks which are approached with a particular objective in mind. How then, should the the sellers be configured in order for the system to stabilise at a particular desired outcome? More specifically, what seller valuation and cost values should be chosen in order for a system containing *spread buyers* to evolve to such a desired outcome allocation?

The method for achieving this is essentially the reverse of the above process. Firstly, using the desired load vector as a starting point, the prices at equilibrium which give rise to the desired loads can be calculated. Secondly, suitable values for the sellers' preference weight and cost parameters are chosen to give rise to equilibrium at the required prices. In this section, this process is explored also by means of some examples; firstly in a baseline case, and then for a more complex objective.

5.4.1 A Baseline Example

Suppose initially that the objective is indeed to achieve an evenly balanced load between two providers in the presence of a population of homogeneous spread buyers. This is equivalent to the assumed objective in chapters 3 and 4. In this instance, the desired outcome may be expressed as follows:

$$\vec{L_S} = \langle \frac{1}{2}, \frac{1}{2} \rangle \,. \tag{5.19}$$

Recall that these values are normalised by $|B|$. However, since the population of buyers is homogeneous, for each seller, $s_i \in S$ it must be that

$$\frac{\sum_{j=1}^{B} q_{ij}}{|B|} = q_{ik}\,, \forall b_k \in B\,, \tag{5.20}$$

where q_{ik} is the quantity of π bought by buyer b_k from seller s_i. Therefore, for any seller s_i, the non-normalised load $l_{s_i} = 0.5 \times |B|$ if and only if $q_{ij} = 0.5$, for all $b_j \in B$. In the case of homogeneous buyers, the required equilibrium prices may therefore be calculated as if there were a single buyer, which for the sake of consistency is denoted by b_j.

Now, given the buyer decision function for ***spread buyers*** (from equation 5.5 above), we know that we require

$$\frac{(v^{\pi}_{b_j} - p^{\pi}_{s_1})}{(v^{\pi}_{b_j} - p^{\pi}_{s_1}) + (v^{\pi}_{b_j} - p^{\pi}_{s_2})} = 0.5\,, \tag{5.21}$$

which simplifies to

$$p^{\pi}_{s_1} = p^{\pi}_{s_2}\,. \tag{5.22}$$

In other words, an evenly balanced load will be the outcome allocation when both sellers quote the same price for π at equilibrium, as we already know. We also know from chapter 3, though it can be easily determined, that this equilibrium may be achieved with zero cost values, $c^{\pi}_{s_1}$, $c^{\pi}_{s_2}$ and when $w_{s_1} = w_{s_2} = 1$.

5.4.2 A More Complex Example

Now let us consider a more complex desired outcome allocation, $\vec{L_S} = \langle \frac{2}{3}, \frac{1}{3} \rangle$. We wish for s_1 to provide twice the load of s_2.

Following the method described in section 5.4.1 above, for a given homogeneous buyer valuation, $v^{\pi}_{b_j}$, the required relationship between $p^{\pi}_{s_1}$ and $p^{\pi}_{s_2}$ at equilibrium may be calculated, in order to achieve the desired outcome:

$$\frac{(v^{\pi}_{b_j} - p^{\pi}_{s_1})}{(v^{\pi}_{b_j} - p^{\pi}_{s_1}) + (v^{\pi}_{b_j} - p^{\pi}_{s_2})} = \frac{2}{3}\,, \tag{5.23}$$

which simplifies to

$$p_{s_2}^{\pi} = \frac{v_{b_j}^{\pi} + p_{s_1}^{\pi}}{2} \,. \tag{5.24}$$

In order to achieve the desired outcome allocation, the prices $p_{s_1}^{\pi}$ and $p_{s_2}^{\pi}$ must conform to this relationship at equilibrium.

The question with which we are faced now becomes what valuation and cost values, $c_{s_1}^{\pi}$, $c_{s_2}^{\pi}$, w_{s_1} and w_{s_2}, can be chosen in order to satisfy this constraint. In order to begin to answer this, we firstly take a look at how the equilibrium prices vary with respect to the differences in seller valuations.

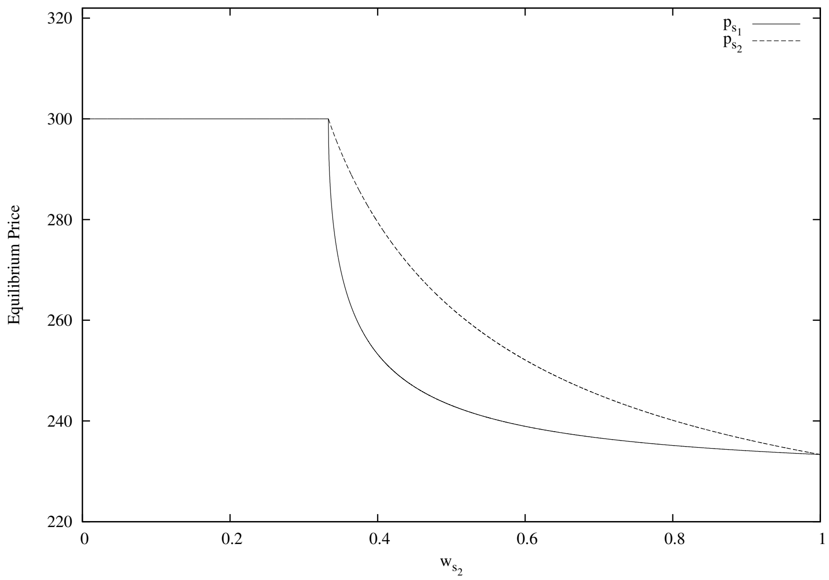

Figure 5.1: Predicted equilibrium prices with respect to seller valuations. $w_{s_1} = 1.0$, w_{s_2} varies along the x axis. First presented in [2].

Figure 5.1 shows two distinct regions within the space of heterogeneous seller valuations. Firstly, note the region when $w_{s_2} < \frac{1}{3}$, where s_2 has no equilibrium price. This may be explained as follows. Since $P_{s_i} = \sum_{j=1}^{|B|} q_{ij}(w_{s_i} p_{s_i}^{\pi} - c_{s_i}^{\pi})$, and $l_{s_i} = \sum_{j=1}^{|B|} q_{ij}$, in order to achieve any positive load l_{s_i}, with a non-negative payoff P_{s_i}, it follows that we must have

$$w_{s_i} \geq \frac{c_{s_i}^{\pi}}{p_{s_i}^{\pi}} \,. \tag{5.25}$$

Now, the maximum possible price $p^{\pi}_{s_i}$ which would lead to any transaction occurring will be the buyers' reserve price, or valuation, $v^{\pi}_{b_j}$. Since, given a fixed cost $c^{\pi}_{s_i}$ this is the price which places a lower bound on the constraint on w_{s_i}, it follows that there exists no price for which the seller s_i has a non-negative payoff when $w_{s_i} < \frac{c^{\pi}_{s_i}}{v^{\pi}_{b_j}}$. In the example illustrated above, this is indeed $\frac{1}{3}$. Below this weight value no transactions will take place, since they would require either buyer or seller to accept a negative payoff.

However, given that s_2 is unable to trade when $w_{s_i} < \frac{1}{3}$ in the above example, in this region s_1 exists in a monopoly. As should be expected, it is therefore able to charge the maximum acceptable price for the buyers, (i.e. $p^{\pi}_{s_1} = v^{\pi}_{b_j}$) and be confident of securing a normalised load of 1.

In the competitive region of the weight space however, the equilibrium prices follow a more complex pattern. Recall that in order to achieve the desired resource allocation, we require a ratio between the sellers' prices as described in equation 5.24. In the example, when $v^{\pi}_{b_j} = 300\,, \forall b_j \in B$, we must require that

$$p^{\pi}_{s_2} = 150 + \frac{p^{\pi}_{s_1}}{2}\,. \tag{5.26}$$

In the example, this occurs when $w_{s_1} = 1.0$ and $w_{s_2} = 0.42$. Therefore, by fixing values for $c^{\pi}_{s_1}$ and $c^{\pi}_{s_2}$, weight values may be identified which achieve the desired outcome allocation of $l_{s_1} = \frac{2}{3} \times |B|$, $l_{s_2} = \frac{1}{3} \times |B|$. Normalised by $|B|$, we therefore have of course $\vec{L_S} = \langle \frac{2}{3}, \frac{1}{3} \rangle$.

5.5 Achieving Desired Resource Allocations

So far this chapter has followed the approach described in chapter 3, focusing largely on predicted outcomes when sellers are able to calculate a best response to the current state of the market. As chapter 4 demonstrated, this may not always provide an accurate representation of the behaviour of a system under the assumption that only private information is available to agents. Therefore, similarly to chapter 4 this section focuses on demonstrating the behaviour of the model in simulation but for the heterogeneous scenario introduced in this chapter.

As with chapter 4, competition between sellers is driven by the co-evolution of their respective *evolutionary market agents*; their evolutionary fitness being analogous to their payoff from interactions in the market, P_{s_i} for seller $s_i \in S$. An *evolutionary market agent*, operating on behalf of a particular resource providing node, has the self-interested objective of maximising its fitness and hence payoff. Section 4.3 described the *evolutionary market agent* algorithm in detail.

5.5.1 Baseline Examples

Firstly, the predicted result in section 5.3.1 is also observed in simulation. Figure 5.2 shows the normalised outcome loads for this two node example scenario. The predicted allocation is quickly achieved, and remains stable.

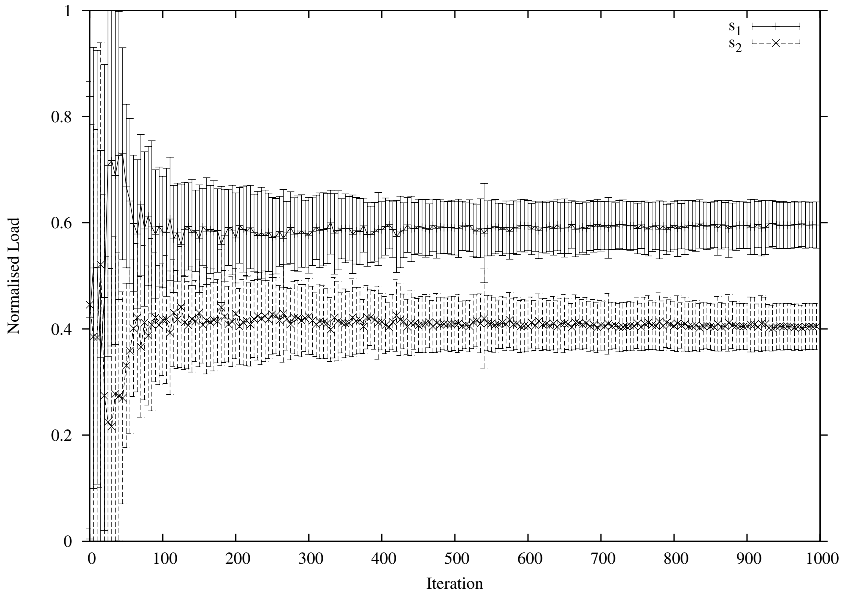

Figure 5.2: Evolution of the loads on s_1 and s_2, where $w_{s_1} = 1.0$ and $w_{s_2} = 0.5$. The resource providers converge to loads of 0.6 and 0.4 respectively. Mean and standard deviation over 30 independent runs. First presented in [2].

Secondly, the case from section 5.4.2 is investigated in simulation, where the desired outcome allocation is $\vec{L_S} = \langle \frac{2}{3}, \frac{1}{3} \rangle$. Figure 5.3 shows the evolution of the load on each resource providing node over time, where $w_{s_1} = 1.0$ and $w_{s_2} = 0.42$ for 30 independent runs; mean and standard deviation are shown. The system achieves the desired allocation, where s_1

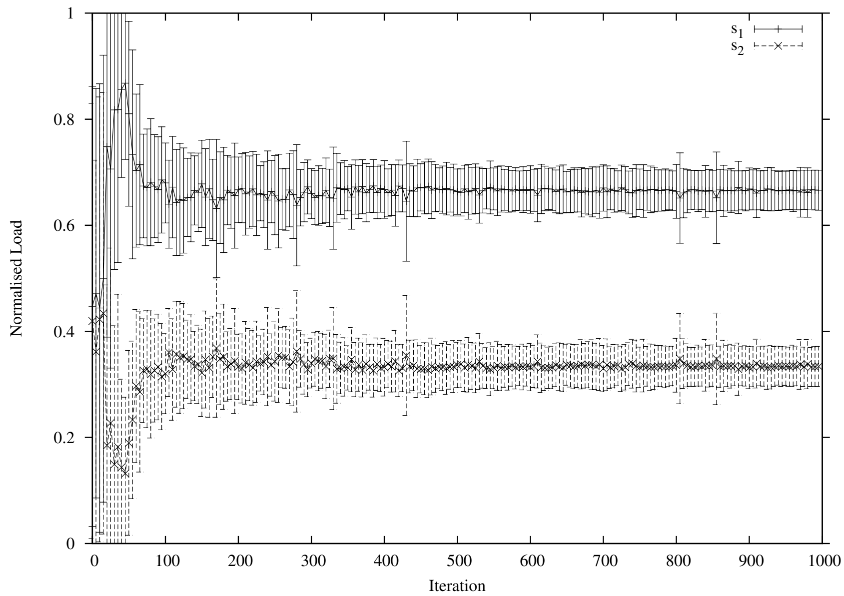

Figure 5.3: Evolution of the loads on s_1 and s_2, where $w_{s_1} = 1.0$ and $w_{s_2} = 0.42$. The resource providers converge to loads of $\frac{2}{3}$ and $\frac{1}{3}$ respectively. Mean and standard deviation over 30 independent runs. First presented in [2].

is providing twice the level of load as s_2, and clearly, as with the performance of the system in the homogeneous case, a rough approximation to the predicted outcome is achieved quickly, while the final allocation remains stable.

5.5.2 A More Complex Example

The more complex four node case described in section 5.3.2 is now examined, in the presence of 100 *spread buyers*. Table 5.1 compares the predicted and experimental outcomes for this example. The experimental results show the state of the system at iteration 1000, for 30 independent runs.

Clearly, the co-evolutionary approach is able to achieve results very close to the predicted outcome, with a high degree of reliability. It is however, worth noting that the unusually high standard deviation of s_4 is due to an outlier in the data obtained. Since the data represents a snapshot of the system at iteration 1000, and given the continued Gaussian mutation of prices about the equilibrium, occasional temporary outliers are to be expected.

Seller	Predicted Price	Experimental Price	Standard Deviation
s_1	214.53	214.79	6.41
s_2	214.53	217.68	6.10
s_3	233.25	233.26	5.98
s_4	200.73	201.62	10.34

Seller	Predicted Load	Experimental Load	Standard Deviation
s_1	0.2536	0.2503	0.0161
s_2	0.2536	0.2535	0.0156
s_3	0.1981	0.1999	0.0133
s_4	0.2946	0.2967	0.0220

Table 5.1: Comparison of predicted and experimental outcome allocations. Experimental results calculated over 30 independent runs. First presented in [3].

5.6 Exploring Equilibria

Finally, section 5.4.2 and particularly figure 5.1 above showed how equilibrium prices vary with respect to differences in seller valuations. But how does this translate into outcome resource allocations? Figure 5.4 illustrates for the two node example studied in that section, the corresponding predicted outcome loads between s_1 and s_2, normalised by the size of the buyer population.

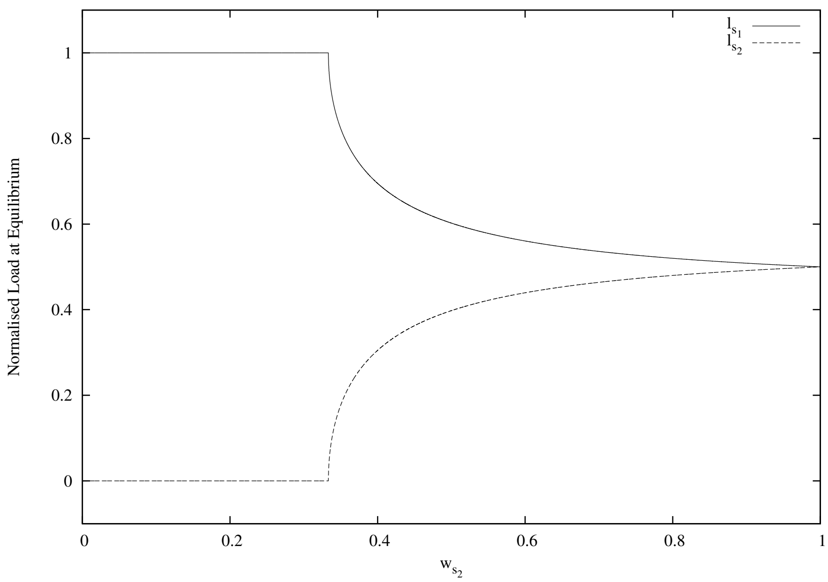

Figure 5.4: Predicted outcome loads with respect to seller valuations. $w_{s_1} = 1.0$, w_{s_2} varies along the x axis. First presented in [2].

Here, the monopolistic region described in section 5.4.2 is clearly visible, when $w_{s_2} < \frac{1}{3}$, as $l_{s_1} = 1$ and $l_{s_2} = 0$. The relationship between the valuation differences and the outcome allocation in the competitive region is however more complex. This is due to the non-linearity of the spread buyer decision function being employed here, and the shape of the curves in this region are dependent upon this behaviour.

This relationship is also observed in simulation, when sellers have access to only private information. Figure 5.5 shows experimental simulation results compared with the predictions in figure 5.4. Mean and standard deviation results are shown, in each case at the allocation reached after 1000 iterations.

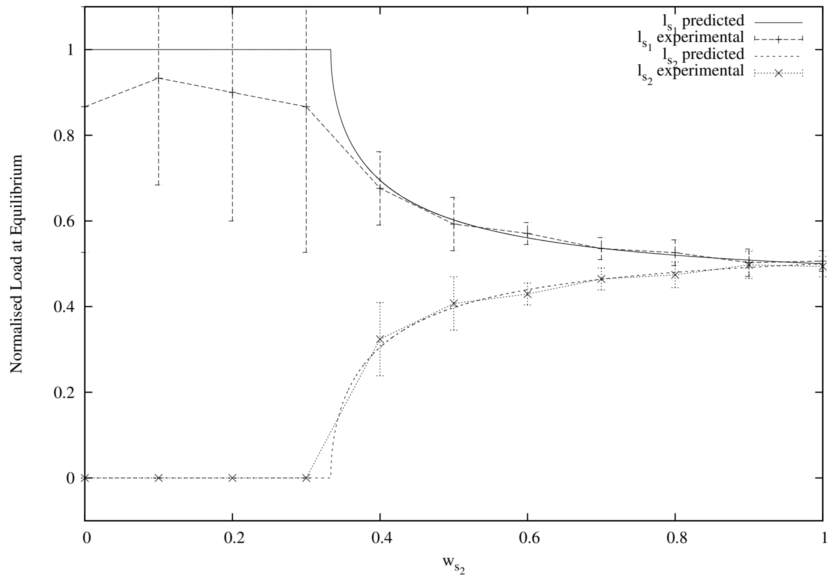

Figure 5.5: Comparison of predicted and experimental outcome loads with respect to seller valuations. $w_{s_1} = 1.0$, w_{s_2} varies along the x axis. Experimental results are mean and standard deviation over 30 independent runs. First presented in [2].

Firstly, note that in the competitive region of the valuation space the experimental results follow the predicted outcome loads to a high degree of accuracy. The stochastic nature of the *evolutionary market agents*' algorithms ensures that there remains some variation at equilibrium, due to ongoing mutation.

During the monopolistic phase however, when $w_{s_2} < \frac{1}{3}$, not only do the experimental results for l_{s_1} only loosely follow their predicted path, but also there is a high standard deviation between the simulation runs. However, further observation reveals that for the experimental results in this region, $l_{s_1} + l_{s_2}$ is often less than 1, meaning that the users' desired provision is not being met in all cases. This is in fact also due to ongoing mutation in s_1's population at equilibrium.

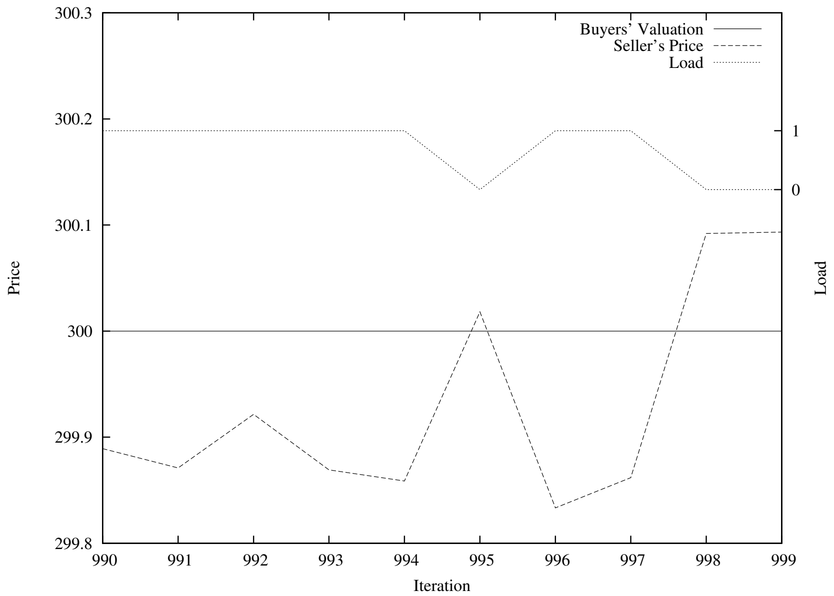

Figure 5.6: Evolution of price in a monopoly. Mutation of the seller's price about the buyers' valuation leads to instability. Results from a typical run when $w_{s_1} = 1.0$ and $w_{s_2} = 0.1$. First presented in [2].

The results in figure 5.6, from a typical run when $w_{s_1} = 1.0$ and $w_{s_2} = 0.1$, illustrate this effect. As the price mutates, there is a probability that it will increase above the buyers' maximum acceptable price. When it does so, no transactions occur, as the seller, s_1 in this case, has priced itself out of the market. The result is temporarily both a load and payoff of zero. Of course, this individual will quickly be selected out, and the price will return to a lower value, however the presence of this step function in the fitness landscape of the *evolutionary market agent* means that this behaviour will continue. Nevertheless, this is not of great concern, since the price's proximity to a step in the fitness landscape only occurs in monopoly, and will not affect regular operation of the algorithm.

5.7 Conclusions and Discussion

This chapter has considered decentralised computational resource allocation when resource providing nodes have heterogeneous capabilities. In this scenario, sellers are considered to value their resources differently, when compared with a common notion of money.

The market-based approach in introduced in chapter 3 was extended and generalised in order to take account of these concepts. It was demonstrated that by accounting for this heterogeneity of seller valuations, stable, desirable uneven outcome resource allocations can be obtained. As with the homogeneous case, this is achieved despite a lack of central control or cooperation between nodes; only self-interested payoff maximising behaviour is relied upon. The possibility of complex buyer decision functions means that this is not a trivial task, since as a result there may not be a straightforward mapping between valuations and the outcome allocation.

Further to this, a method was described for determining suitable parameters for the resource providing nodes, given a desired outcome resource allocation and buyer behaviour. The ability of the approach to achieve a desired uneven allocation is achieved transparently to resource users, as no modification to the buyers' behaviour is required.

It is therefore possible to predict the outcome resource allocation for given sets of buyers and sellers, or conversely to parametrise them in order to achieve a desired outcome. Sellers taking account of the cost of provision, and their valuation of money when compared with this cost, leads to a resource allocation which reflects the varied ability of resource providers. Furthermore, simulations making use of *evolutionary market agents* demonstrated that the predicted outcomes for *spread buyers* are also obtained when sellers have only private information available, as in the homogeneous case.

This chapter's contributions are therefore the following:

- An extension of the posted offer market-based approach described in chapter 3 to consider heterogeneous resource providing node capabilities, by means of additional cost and valuation parameters.

- A game theoretic analysis of how the extended model can obtain additional stable, predictable, uneven outcome resource allocations at equilibrium.
- Experimental evidence which demonstrates that these outcomes can be achieved under the assumption of only private information, when sellers use *evolutionary market agents.*
- A method of determining suitable seller cost and valuation parameters in order to elicit the additional desired outcome resource allocations.

CHAPTER 6

USING MULTIPLE ATTRIBUTES TO ATTAIN DESIRED QUALITIES OF SERVICE

A question of need is a question of taste.

Neil Tennant and Chris Lowe

This chapter builds upon the previously investigated retail-inspired approach, in which resource providing nodes compete on price to sell their resources in a posted offer market mechanism. The extension presented in this chapter allows for resources to be described over additional quality attributes, in addition to simply their price. In this extended multiple attribute model, resource users may have different preferences and constraints over the attributes, and as such form niches in the market. Sellers therefore may target particular market segments.

A game theoretic analysis of sellers' behaviour in the multiple attribute version of the model is presented, and results compared for when buyers act as *bargain hunters* and *spread buyers*. It is demonstrated that in the presence of *bargain hunters*, niches in the market emerge, but *myoptimal* selling agents become trapped in an unending cyclical competitive dynamic, similar to the price wars described by Kephart et al. [9]. When buyers adopt *spread buyer* behaviour, the cyclical behaviour is no longer present, and sellers instead converge to an equilibrium state, providing their resources with price and quality attributes appropriate to the buyer population. Qualitative differences are highlighted between two attribute and three attribute scenarios; in the latter case market segmentation is no longer observed with *spread buyers*.

Section 6.1 begins the chapter by motivating the extension to the original model presented in chapter 3. In section 6.2 the multiple attribute extension is described, including a discussion of agent behaviour and potential utility models. In section 6.3 the multiple attribute payoff landscape is introduced and section 6.4 investigates competitive seller dynamics in the model with a homogeneous buyer population. The segmentation of selling agents into niches is illustrated in sections 6.5 and 6.6 through a series of examples. Throughout, the effects of both *bargain hunters* and *spread buyers* are considered. Finally, section 6.7 presents a discussion on what can be learnt from this analysis.

6.1 Motivation

So far, we have seen how posted offer markets may be used to allocate instances of a resource amongst a set of agents, using evolutionary algorithms as an adaptive pricing mechanism. Hitherto however, the resource has been assumed to be homogeneous in type, a unit of resource from one seller is perfectly substitutable for that from another. The competitive element of the systems investigated has been the price charged by each seller per unit of the resource in question.

In many scenarios this will be sufficient, since many computational resources are equivalent regardless of their source. For example, in a homogeneous compute cluster environment, it is unlikely that a user will have a preference over which available node is used to perform the computation, since the end result is the same. Similarly, a web service providing accurate time may be offered by any number of functionally correct servers, and from the user's perspective, which one is chosen will not matter.

However in other scenarios, users of computational resources will vary in their non-functional requirements. For example, Zeng et al. [118] highlight that users of web services will have different preferences and constraints regarding quality of service (QoS) attributes, such as execution time and reliability. They argue that mechanisms for the allocation of resources in such systems should be designed to take account of such a heterogeneity of constraints on and preferences between non-functional requirements.

6.2 Extending the Market Model with Multiple Attributes

The model used here is an extension of that described in chapter 3, the retail-inspired posted offer market. Chapters 4 and 5 showed that this can be a useful mechanism for achieving a balanced load as well as other stable resource allocations in decentralised computational systems. In the original model, resources are described purely in terms of their price, and are assumed otherwise to be perfectly substitutable.

Using the original notation, a resource π is offered by a seller s_i for the price $p^{\pi}_{s_i}$. If π

may instead be described over a number of attributes π_1, π_2, ... π_n, then the offered price $p^{\pi}_{s_i}$ may be replaced by a vector of offered attribute values, $\vec{x}^{\pi}_{s_i}$. For example:

$$\vec{x}^{\pi}_{s_i} = \langle x^{\pi_1}_{s_i}, x^{\pi_2}_{s_i} \rangle \tag{6.1}$$

6.2.1 Agent Behaviour

As in the original model, buying agents are faced with a decision problem: given the offers available, which one(s) should be chosen? Sellers are faced with a similar decision: what offer will provide the maximum payoff given the buyers and competitors? In order to decide this, the agent is endowed with a utility model which reflects the decision making priorities of its owner: the user or supplier of the resource.

To a certain extent, it does not matter what this model is; the agent may treat its utility function as a black box, into which it places an offer, and from which it receives a utility value for that offer. In this way, offers may be compared against each other both in terms of preference and the magnitude of the preference.

However, in order to aid analysis at this stage, a common utility model from multi-attribute utility theory (MAUT) [49], a branch of decision theory concerned with multiple criteria, is used. Its aim is to analyse and compare the attractiveness of a set of alternatives, by algebraically calculating a single conjoint utility value for each alternative. This approach has long been useful in real world decision making [119, 120, 121], where an expert's time is valuable or limited, and an analytical model may be used as an aid to the decision process. Similar reasoning applies here, where for example, the user of a resource may not be willing or able to assess each and every offer for its attractiveness themself, and an agent in possession of a model may save considerable time and effort.

The translation of human or business needs and preferences into utility functions and weights for their respective agents is a large and complex task, and this thesis does not pretend to contribute to that field. Nevertheless, the commonly used additive utility function of MAUT [49], which is relevant to many real world decision processes is that which is used in this chapter.

It is important however to note the assumptions relied upon by MAUT, particularly that of ***utility independence*** of attributes. Briefly, if the utility function of attribute x does not shift strategically as attribute y varies, then x is ***utility independent*** of y. In other words, for x to be ***utility independent*** of y, all the utility functions for x, defined as y varies, must be positive linear transformations of each other [49].

6.2.2 Additive Utility Model

As discussed, agents' utility functions may be constructed in a variety of ways in order to represent the underlying desirability of different combinations of attributes. One of the simplest utility models in MAUT is the additive utility function [49]. The additive utility model assumes a refinement of utility independence: ***additive independence*** of the attributes. For attribute x to be ***additively independent*** of attribute y, the utility contribution of x must not be dependent on the value of y. If this assumption is true for all attributes, then the additive model may be used. For a more formal definition of ***additive independence***, see Keeney and Raiffa [49].

In the additive model the decision maker calculates a weighted sum of the utility associated with each attribute. An attribute's weight represents the importance of that attribute to the decision maker, and gives its contribution to the overall utility of the alternative. The outcome of the function is a conjoint value for each alternative, which may be used to determine the overall relative preference. Describing an agent a's consideration of an offer by seller s_i, the agent's utility would be as follows:

$$u_a^\pi(\vec{x}_{s_i}^\pi) = \sum_{j=1}^{n} k_a^{\pi_j} u_a^j(x_{s_i}^{\pi_j}) \tag{6.2}$$

The attribute's utility function u_a^j determines the unweighted contribution of the jth attribute to agent a's utility. The attribute's weight or scaling constant, $k_a^{\pi_j}$ represents the importance of the attribute to the decision maker, agent a.

The additive utility model is designed such that attribute utility contributions normalised to between 0 and 1 yield an outcome conjoint utility also between 0 and 1. Therefore the scaling constants are also normalised, such that $\forall k_a^{\pi_j}, 0 \leq k_a^{\pi_j} \leq 1$ and $\sum_{j=1}^{n} k_a^{\pi_j} = 1$.

6.2.3 A Simple Example

The decision process may be illustrated by means of a simple example. Consider a buyer b faced with the following offers from sellers s_1 and s_2:

$$\vec{x}_{s_1}^{\pi} = \langle 600, 300 \rangle$$

$$\vec{x}_{s_2}^{\pi} = \langle 800, 500 \rangle$$

where a high value for the first attribute is undesirable to the buyer, but a high value for the second attribute is desirable. The buyer agent's attribute utility functions will of course determine the final utility value, but in this example assume that they are linear functions which normalise from between some minimum value 0 and maximum value 1000, and in the case of the first attribute, obtains the inverse. This may result in a conjoint utility of, for example,

$$u_b^{\pi}(\vec{x}_{s_1}^{\pi}) = -0.6k_b^{\pi_1} + 0.3k_b^{\pi_2}$$

$$u_b^{\pi}(\vec{x}_{s_2}^{\pi}) = -0.8k_b^{\pi_1} + 0.5k_b^{\pi_2}$$

Since $u_b^{\pi}(\vec{x}_{s_1}^{\pi}) > 0$ if and only if $k_b^{\pi_2} > 2k_b^{\pi_1}$, the offer from s_1 is only acceptable to b if the second attribute is twice as important as the first. For s_2's offer, $u_b^{\pi}(\vec{x}_{s_2}^{\pi}) > 0$ if and only if $k_b^{\pi_2} > 1.6k_b^{\pi_1}$, which is somewhat less demanding.

Note that though this example is described in terms of a buyer agent, the same approach may be applied to sellers, when comparing the utilities of two or more offers.

6.2.4 Budget Constraints

The utility functions discussed above describe how an agent may order and compare offers described over multiple attributes. However, they do not consider that an offer, though attractive, may be unattainable. A simple example of this is an offer which, from a buyer's perspective, provides a highly attractive resource at a reasonable but unaffordable cost. In this scenario, an undesirable attribute such as price, may be beyond the buyer's budget constraint.

Similar constraints may exist on desirable attributes; a buying agent may be instructed not to accept an offer where for example, a security attribute is below a certain threshold. Similarly, a seller agent may be physically unable to provide a resource with a latency below a certain limit value. Regardless of its utility, an offer where any attribute is beyond the budget constraint of the decision maker is therefore considered unacceptable and it is rejected.

In this thesis, lower and upper constraint values are denoted by a line under or over the attribute value respectively. For example, $\underline{x}_s^{\pi_1} = 100.00$ and $\overline{x}_b^{\pi_1} = 300$ indicates that seller s will not accept an offer for which x^{π_1} is below 100.00, and buyer b will not accept an offer for which x^{π_1} is above 300.

6.3 Baseline Scenario: Seller Unit Payoff

In the extended multiple attribute model, two attributes may be used to represent notions of price and quality. Let us examine how buyers and sellers will interact in this new scenario. Figure 6.1 shows the payoff landscape for a seller selling one unit of the resource π. The seller uses the additive utility model, as described in section 6.2.2, with linear attribute utility functions, such that positive values of π_1 lead to positive utility values and positive values of π_2 lead to negative utility values. Its preference weights are $0.7, 0.3$ for attributes π_1 and π_2 respectively. Higher payoff is indicated by the lightness of the blue area, while the region in which offers result in negative payoff is indicated in red, with the darker the red meaning the greater the loss of payoff per transaction.

Unsurprisingly, if the sale is guaranteed, the seller can obtain its maximum payoff by

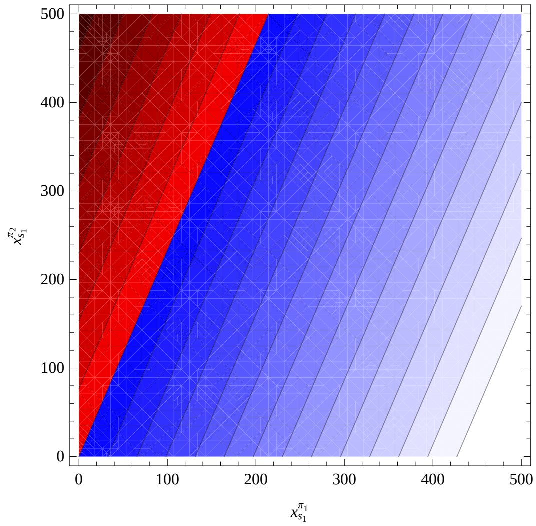

Figure 6.1: Two attribute example payoff landscape for seller s_1, for the sale of a single unit of resource π.

maximising $x_{s_1}^{\pi_1}$ (notionally price) and minimising $x_{s_1}^{\pi_2}$ (notionally quality). Note that the structure of this landscape remains the same as the quantity of π sold increases; the payoff merely scales uniformly with the quantity.

6.4 Homogeneous Buyers

Suppose now that there is a homogeneous set of buyers, B, to whom the resource is offered. To avoid abusing notation too much, this analysis does not refer to the whole set of buyers, but instead to a single buyer $b \in B$. However, since all the buyers in B are the same, the analysis of b may be seen as representative of B as a whole. Each buyer has an upper budget constraint, as described in section 6.2.4, in terms of attribute π_1.

6.4.1 Monopolistic Scenario

Firstly, let us consider the payoff available to the seller in the monopolistic case, that when there is no competition from any other sellers. Figure 6.2 shows the payoff landscape in the presence of a homogeneous population of buyers, where the buyers' budget constraint $\overline{x}_b^{\pi_1}$ is set at 200.00.

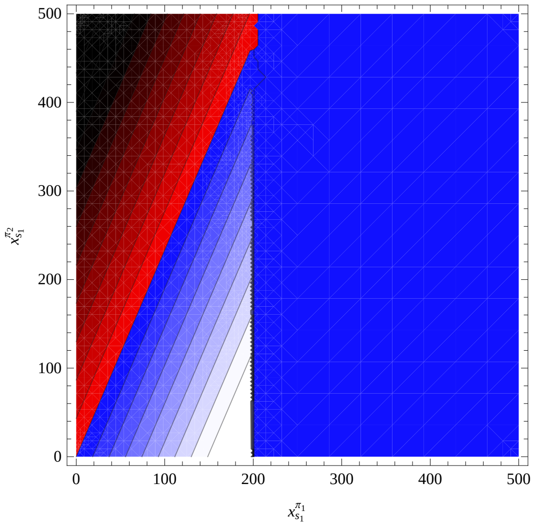

Figure 6.2: Two attribute example payoff landscape for seller s_1, for the sale of a single unit of resource π, where the buyer has a upper constraint of 200.00 on attribute π_1.

Here the optimal offer from seller s_1's perspective is still simply to minimise $x_{s_1}^{\pi_2}$ and to maximise $x_{s_1}^{\pi_1}$ as much as possible within the buyers' budget constraint. Any offer for which $x_{s_1}^{\pi_1}$ is greater than the budget constraint of 200.00 will be rejected, leading to a payoff of zero for the seller.

Now consider also that the buyer makes use of an additive utility model with linear attribute utility functions, as described in section 6.2.3. Figure 6.3 shows the seller's payoff landscape with such a buyer whose preference weights $k_b^{\pi_1}$ and $k_b^{\pi_2}$ are 0.3 and 0.7 respectively.

Note also that due to the seller's own preference weights, if $x_{s_1}^{\pi_1}$ is sufficiently low in relation to $x_{s_1}^{\pi_2}$, such that $k_{s_1}^{\pi_1} u_{s_1}^{\pi_1}(x_{s_1}^{\pi_1}) < k_{s_1}^{\pi_2} u_{s_1}^{\pi_2}(x_{s_1}^{\pi_2})$, then the seller's payoff drops below zero. This is indicated by the line separating the blue and red regions of the figure.

Therefore, there exists a feasible region within the attribute space, bounded by the triangle $(0,0)$, $(\overline{x}_b^{\pi_1}, \frac{k_b^{\pi_1}\overline{x}_b^{\pi_1}}{k_b^{\pi_2}})$, $(\overline{x}_b^{\pi_1}, \frac{k_{s_1}\overline{x}_b^{\pi_1}}{k_b^{\pi_2}})$. Given our assumption of only voluntary interaction between buyers and sellers, offers outside this region will be rejected either by the seller or the buyer, since the resulting payoff will be below zero for at least one of them.

Indeed, in this monopolistic example, the optimal place for our seller to position itself is at the buyer's budget constraint in terms of attribute π_1, such that $x_{s_1}^{\pi_1} = 200.00$ in the

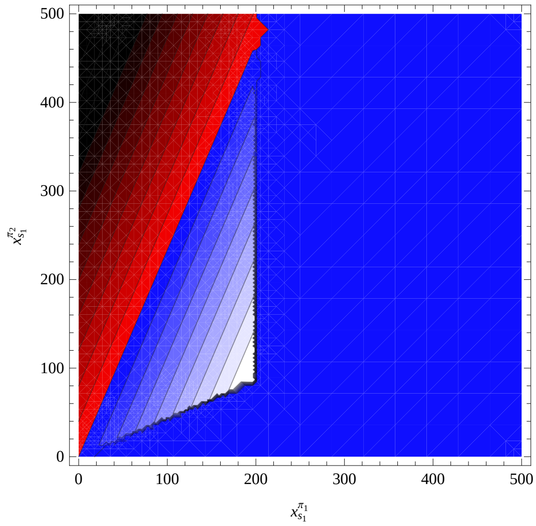

Figure 6.3: Two attribute example payoff landscape for seller s_1, for the sale of a single unit of resource π, where the buyer has a upper constraint of 200.00 on attribute π_1 and preference weights $0.3, 0.7$.

example, and at the corresponding limit value for attribute π_2, $x_{s_1}^{\pi_2} = \frac{0.3*200.00}{0.7} = 85.71$; the optimal offer for the seller is at the lower right corner of the triangle.

6.4.2 Competitive Scenario: Bargain Hunters

Now, if the seller s_1 is in competition with another seller, say s_2, the payoff landscape for each will also depend on the position of their competitor's offer. Once the optimal position in the attribute space, as identified in section 6.4.1 is occupied by a competitor, the seller's payoff landscape has shifted.

Suppose, as in the previous section, s_1 has positioned itself at $200.00, 85.71$. If the competing seller s_2 were to occupy exactly the same position, then the quantity sold to the two sellers would be expected to be equal, and the payoff available would be shared between them; each seller would see their payoff decrease by half.

At this stage it is necessary to consider the behaviour of the buyers to whom the sellers are selling. For consistency, the buyer types introduced in section 3.1.3 are considered here. If the buyers are *bargain hunters*, and therefore purchase everything from the single most attractive seller, then rather than position itself at the optimal monopolistic position in the

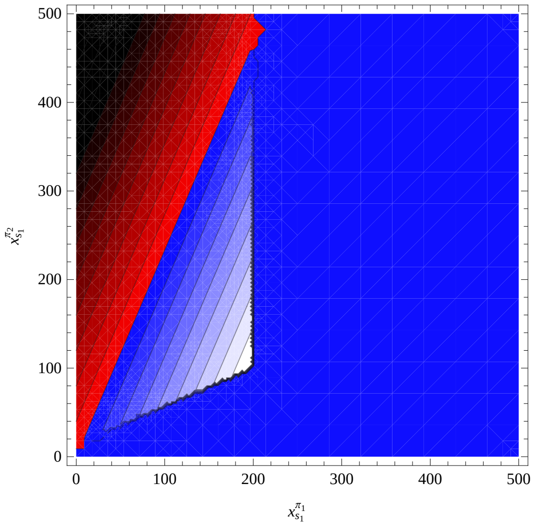

Figure 6.4: Two attribute example payoff landscape for seller s_1, given one competitor offering $200.00, 100.00$ and a homogeneous population of ***bargain hunters***.

attribute space, s_2 can obtain a higher payoff by improving on this offer by a small amount, when viewed from the buyers' perspective. This may involve either reducing $x_{s_2}^{\pi_1}$ (price) or increasing $x_{s_2}^{\pi_2}$ (quality), or both, by some small amount.

The distance from s_1's offer within which a positive difference will be realised to s_2's payoff will depend on the gradient of the payoff landscape around the s_1's position, though in our example, suppose s_2 increases $x_{s_1}^{\pi_2}$ from 85.71 to 100.00, while keeping $x_{s_1}^{\pi_2}$ the same. By doing so, it can expect to earn a greater payoff than half of that available at the monopolistic position.

As a result of this move, the payoff landscape for the original seller, s_1 is also modified. Figure 6.4 shows the resulting payoff landscape for s_1, given that s_2 is offering $200.00, 100.00$. The structure of the landscape is identical to in figure 6.3 except that as a result of s_2's action, the bottom right corner of the triangle bounding the feasible region for s_1 is now at the position of s_2's offer. Any offer which would be less attractive to the buyers than that of s_2 would be rejected.

Therefore, s_1 improves its offer also, and subsequently s_2 is required to do likewise. As a result, the feasible region of the payoff landscape decreases in size, and the peak also drops in height. Figure 6.5 shows the resulting landscape when competition has driven the offers to

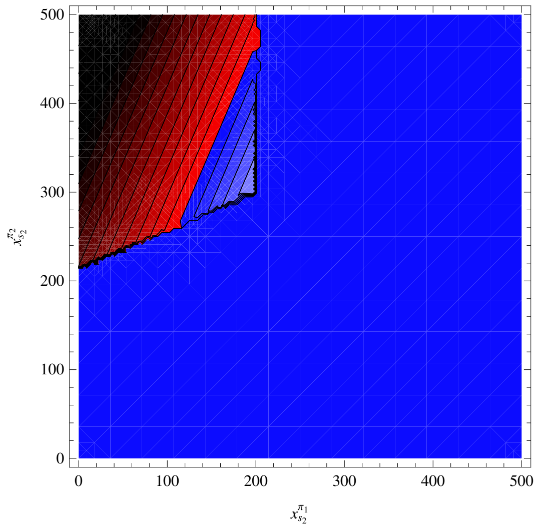

Figure 6.5: Two attribute example payoff landscape for seller s_2, as a result of competition on attribute π_2, with a homogeneous population of ***bargain hunters***.

$200.00, 300.00$. The colour scaling has been kept consistent throughout this series of figures, to indicate that the height of the peak decreases as competition drives $x_{s_1}^{\pi_2}$ and $x_{s_2}^{\pi_2}$ upwards.

This sequence of competitive adaptation and co-adaptation culminates in $x_{s_1}^{\pi_2}$ and $x_{s_2}^{\pi_2}$ reaching the sellers own limit values, determined by their utility functions, at which point both sellers receive no payoff, and have no option to do otherwise.

It is also possible, of course, that rather than increase $x_{s_1}^{\pi_2}$, the seller could respond to competition by decreasing $x_{s_1}^{\pi_1}$. Recall the payoff landscape before competition ensued, depicted in figure 6.3. However, this would in fact be a fruitless move. Figure 6.6 shows the payoff associated with the attribute space from the buyer's perspective. Here, light blue is used to indicate high payoff for the buyer, and dark red to indicate greater losses on the part of the buyer.

It is clear that when faced with a choice between two sellers, one which has responded to competition by increasing x^{π_2} and another by reducing x^{π_1}, that the buyers would prefer the seller with the higher value for x^{π_2}, maintaining x^{π_1} at the budget constraint. For this reason, sellers dropping below the budget constraint for x^{π_1} would be expected to therefore quickly respond by returning to the budget constraint line.

We therefore observe that self-interested competition between the two sellers described

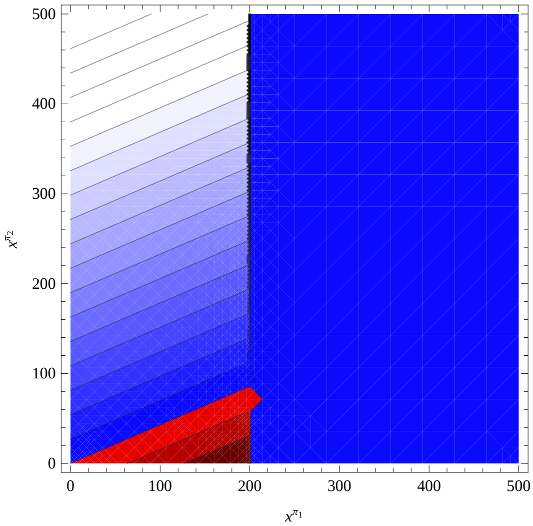

Figure 6.6: Two attribute example payoff landscape from the buyers' perspective. Compare with figure 6.3.

in the example will result in $x_{s_1}^{\pi_1}$ and $x_{s_2}^{\pi_1}$ remaining at the budget constraint, while $x_{s_1}^{\pi_2}$ $x_{s_2}^{\pi_2}$ increase due to competitive pressure, until the resulting payoff for both competing sellers is zero.

The analysis presented here is indeed a two-dimensional version of that applied to the single attribute case, described as a *price war* in section 3.2, and is reminiscent of Bertrand competition, where sellers compete on price alone. The Nash equilibrium here is for both sellers to make offers at the top corner of the triangle bounding the feasible region, which is unaffected by competition and positioned at $(\overline{x}_b^{\pi_1}, \frac{k_{s_1}\overline{x}_b^{\pi_1}}{k_b^{\pi_2}})$. Both sellers will both receive a payoff of zero.

6.4.3 Competitive Scenario: Spread Buyers

As was shown in section 3.2, in the single attribute case *spread buyers* give rise to a somewhat different outcome; the same is true in the multiple attribute case. Since the differences between *bargain hunters* and *spread buyers* are concerned only with how multiple sellers making acceptable offers are treated by the buyers, in the absence of any competition a seller is faced with exactly the same payoff landscape as in figure 6.3 above. When there is only one seller

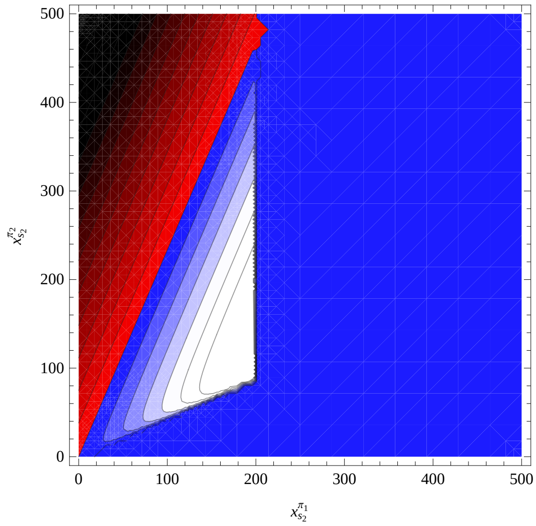

Figure 6.7: Two attribute example payoff landscape for a seller s_2, given one competitor offering $200.00, 85.71$ and a homogeneous population of *spread buyers*.

present, it is of course able to behave as a monopolist, as described in section 6.4.1.

When faced with a choice between two or more competing sellers however, the *spread buyers* will give rise to a different payoff landscape. For ease of comparison, in the following analysis all else remains equal between this and the *bargain hunter* case in section 6.4.2. Suppose that initially seller s_1 offers the monopolistically optimal offer of $200.00, 87.51$, as it was shown to be in section 6.4.1. A competing seller, s_2 is then faced with the landscape shown in figure 6.7, which may be compared with that in figure 6.4.

Here the familiar shape of the payoff function arising from spread buyers, introduced in section 3.2.3 can be seen. However, this time the curve projects horizontally from the budget constraint line. As with *bargain hunters*, there is a clear peak in the landscape for the seller to occupy; in this example the optimal, or best response offer is $200.00, 110.14$. Assuming that s_2 does indeed respond with this, the payoff landscape for its competitor s_1 becomes that depicted in figure 6.8.

As can be seen, both the location and curve of the peak shifts with this response. For s_1 the best response is now $200.00, 159.72$. The equilibrium offer may be found by calculating best responses in an iterative manner, as was demonstrated in the single attribute case in section 3.2.3 and used again in section 5.4. In the example given here, equilibrium is reached

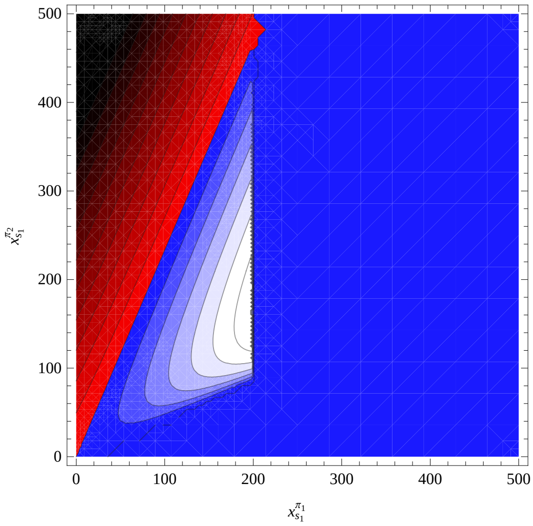

Figure 6.8: Two attribute example payoff landscape for a seller s_1, given one competitor offering $200.00, 110.14$ and a homogeneous population of *spread buyers*.

when both sellers offer $200.00, 212.68$. At this stage, the payoff landscape for both sellers is that in figure 6.9, where the peak in the landscape, the best response is the same as that which both sellers made previously. The system is stable.

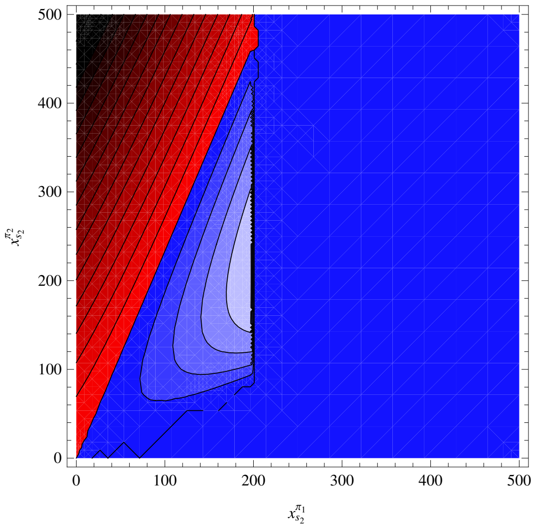

Figure 6.9: Two attribute example payoff landscape for a seller s_2, given one competitor offering the equilibrium offer of $200.00, 212.68$ to a homogeneous population of *spread buyers*.

6.5 Heterogeneous Buyer Constraints

Of course, as discussed in section 6.1, not all buyers can be expected to have the same budget constraints or preferences. A heterogeneous population of buyers will lead to a more complex payoff landscape for the sellers, and result in more complex dynamics between them. In this section, the model is analysed when buyers vary in their budget constraints.

It has long been understood [122] that market segmentation is a likely outcome of markets with heterogeneous buyers, since it becomes advantageous for sellers to adapt to the different needs of the buyers. This is the process by which sellers occupy niches in the market, each niche appropriate to a particular profile of buyer. As Dickson and Ginter [122] describe, the market may be viewed as a set of segments, or sub-markets, which sellers may target.

As this section will demonstrate, the presence of heterogeneous constraints in a population of buyers is sufficient to lead to the presence of market niches even when in all other respects the buyers are homogeneous. Here this is referred to as *constraint-induced niching*.

6.5.1 Monopolistic Scenario

First consider the payoff landscape for a seller in the presence of a population of *bargain hunters*. Here the utility functions and preferences remain as in the example in the previous sections, however half of the buyers have a budget constraint of 200.00 and half of 100.00, in terms of attribute x^{π_1}. This landscape is shown in figure 6.10, which may be compared with the homogeneous case in figure 6.3. Colour scaling has been reset from the previous figures, for clarity's sake.

As can be seen, there are now two peaks in the payoff landscape, each of a similar structure to in the homogeneous case. Indeed, these might be described as two sub-markets in the same attribute space. In this particular case, the two peaks provide equivalent payoff for a monopolistic seller, though this will not generally be true. Figure 6.11 shows two further potential landscapes generated only through changes to the buyers' budget constraints.

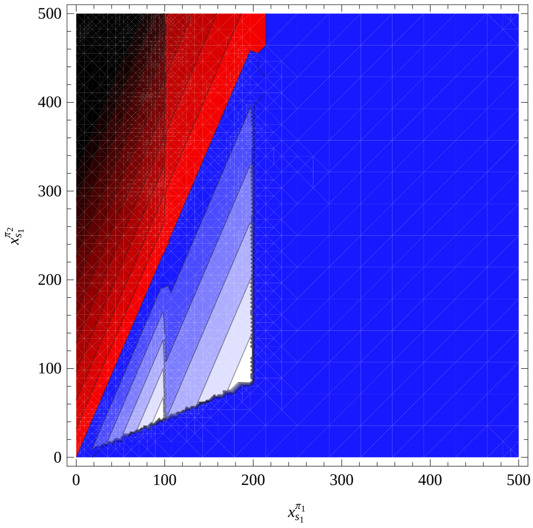

Figure 6.10: Two attribute example payoff landscape for a seller s_1, with a heterogeneous population of buyers. Half of the buyers have a budget constraint of 100.00, half 200.00.

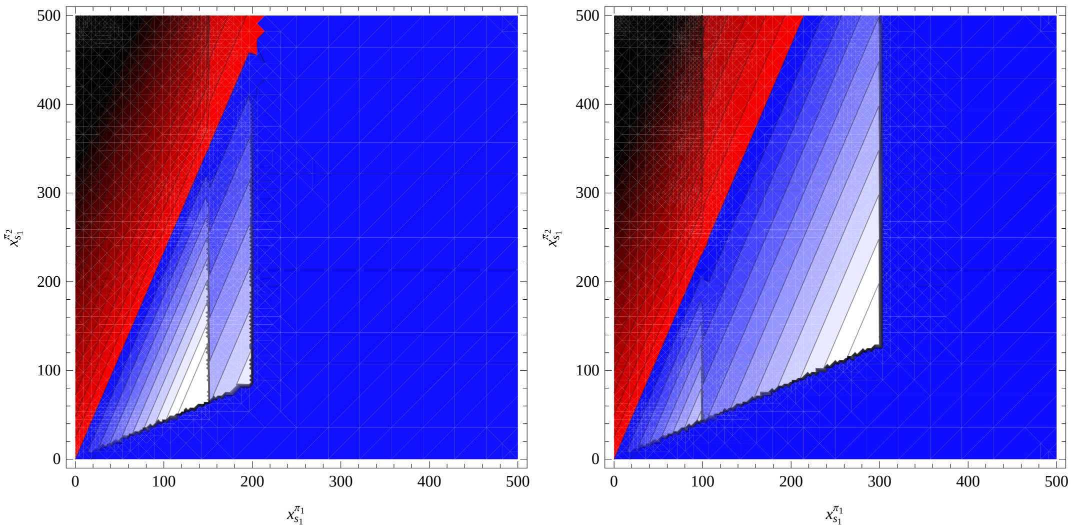

Figure 6.11: Two attribute example payoff landscapes for a seller s_1, with a heterogeneous population of buyers. The buyers' budget constraints are split equally between 150.00 and 200.00 on the left plot and 100.00 and 300.00 on the right one.

6.5.2 Competitive Scenario: Bargain Hunters

As in the homogeneous case, let us now consider the landscape as presented to a seller s_2 with a single competitor. Suppose given the landscape in figure 6.10, s_1 offers $200.00, 87.51$, at the right peak. This has the same effect on this particular segment as was seen in the homogeneous case in section 6.4.2. s_2 is able to increase its payoff by improving its offer over s_1, and which point s_1 responds by improving its offer further. Competition ensues for the provision of the resource to the right sub-market. Figure 6.12 shows the payoff landscape presented to s_2 during this price war, specifically when s_1 is offering $200.00, 150.00$.

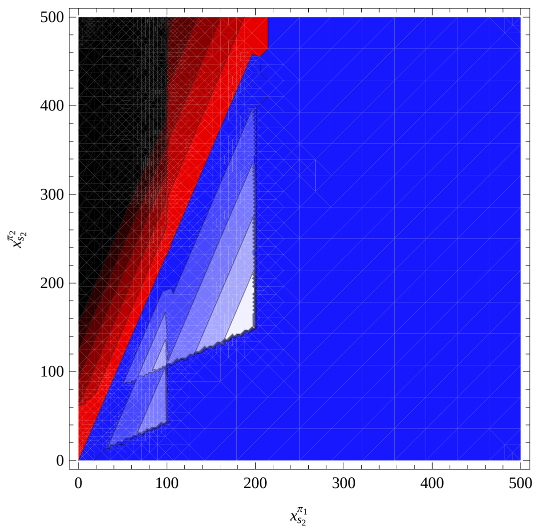

Figure 6.12: Two attribute example payoff landscape for a seller s_2, with a heterogeneous population of *bargain hunters*. Half of the buyers have budget constraints of 100.00, half 200.00. One competing seller offers $200.00, 150.00$.

As can be seen, competition on the right peak actually has the effect not only of bringing down the payoff for those competing for it, as was seen in the homogeneous case, but also of reducing the payoff available for a seller contemplating the lower peak. Here, the seller can of course respond by making an improved offer over its competitor on the right peak, and continuing the escalation as was the only option in the homogeneous example. This will of course ultimately bring down the payoff available at this peak until it reaches zero. However, in the heterogeneous case, the seller can also respond by moving to the lower peak, albeit at a lower payoff, but avoiding the competitive pressure.

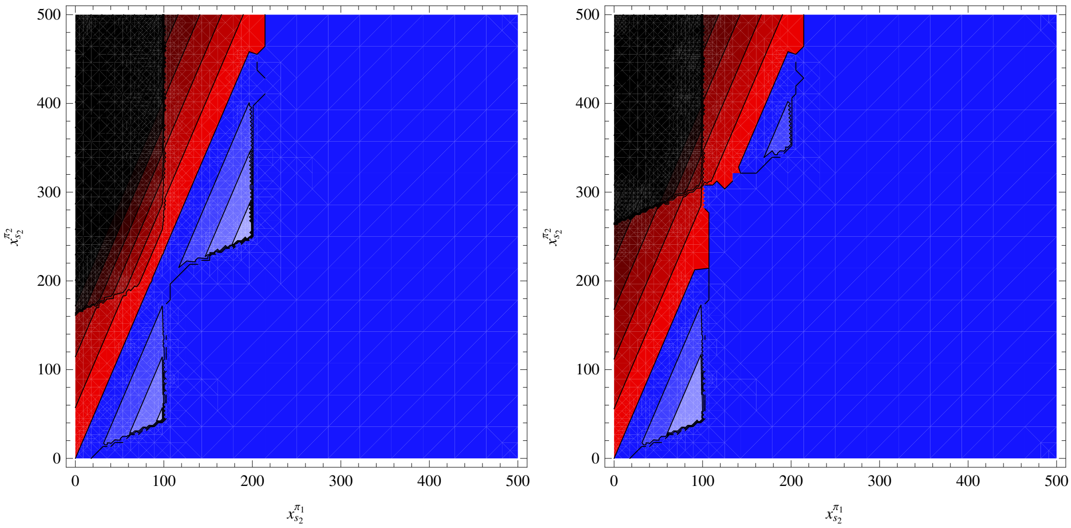

Figure 6.13: Two attribute example payoff landscape illustrating competition between two sellers on a single peak. The payoff available from the buyers being competed for decreases until it is lower than that available at the other market niche. If competition continues, eventually it reaches zero.

Suppose first that competition does continue on the right peak. Figure 6.13 shows how competition reduces the payoff available as the right peak continues to shift towards higher values for x^{π_2}. Firstly, the peaks become roughly equal in height. Subsequently, the higher valued peak reduces close to zero payoff. Of course, unlike in the homogeneous case, once competition has reached the zero payoff state, the system is not stable, as the left peak now clearly offers a higher payoff for any seller wishing to move to it.

As can be seen, once the payoff received by a seller competing for the right peak is sufficiently low, it may increase its payoff by moving its offer to the left peak in order to target the other niche. Figure 6.14 shows the payoff landscape facing seller s_2 once its competitor s_1 has made this jump to the lower valued peak with an offer of $100.00, 40.00$.

As can be clearly seen, immediately the two original peaks are recreated, as the competitive pressure is removed. Of course, subsequent competition on either peak will lead to another escalation of x^{π_2}, until again it appears beneficial to switch to the other peak, and the process begins again. Figure 6.15 illustrates this effect this time on the left peak.

The cycle of competition observed in this model is the same effect as observed by Kephart et al. [9] in their study of a model of a computational market in which agents compete

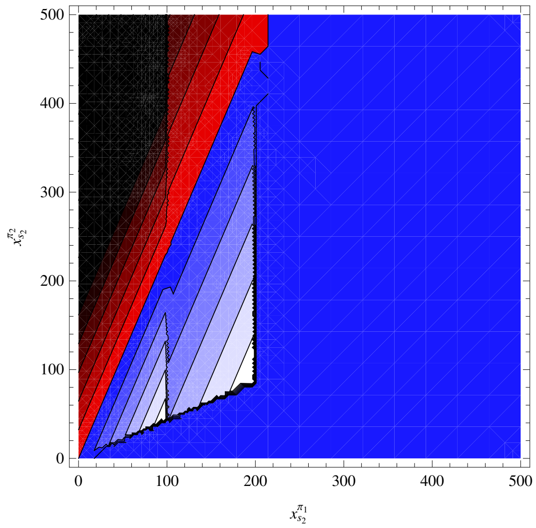

Figure 6.14: Two attribute example payoff landscape for a seller s_2, with a heterogeneous population of *bargain hunters*. Half of the buyers have budget constraints of 100.00, half of 200.00. One competing seller offers $100.00, 40.00$.

to provide relevant informational services to consumers. They show that when sellers are myopically optimal, or *myoptimal* the presence of multiple peaks in their payoff landscape leads to an unending series of price wars, interrupted by abrupt resets. In the example presented in this section, competition is not on price, but instead on another attribute, but the logic behind it remains the same. Kephart et al. argue that these *limit-cycle wars* are qualitatively the same as those observed in a model by Edgeworth with a similar two-peak payoff landscape, despite the causes for the multiple peaks being different.

As has been shown, a seller's decision on which sub-market to target is not straightforward, since the relative heights of the peaks vary with the competitor's offer. Figure 6.16 shows for s_2 the height of each peak, the maximum payoff available by targeting each market segment, as a competitor s_1 varies its offer. On the left, the competitor explores the left peak by fixing $x_{s_1}^{\pi_1}$ at 100.00 and varying $x_{s_1}^{\pi_2}$ (shown on the horizontal axis). The lines show the subsequent maximum payoff obtainable by s_2 in each niche. Similarly, on the right s_1 explores the right peak, fixing $x_{s_1}^{\pi_1}$ at 200.00 and again varying $x_{s_1}^{\pi_2}$. The maximum obtainable payoff for s_2 is again shown for each niche.

The most notable feature of the plots in figure 6.16 is that in the right hand plot the lines

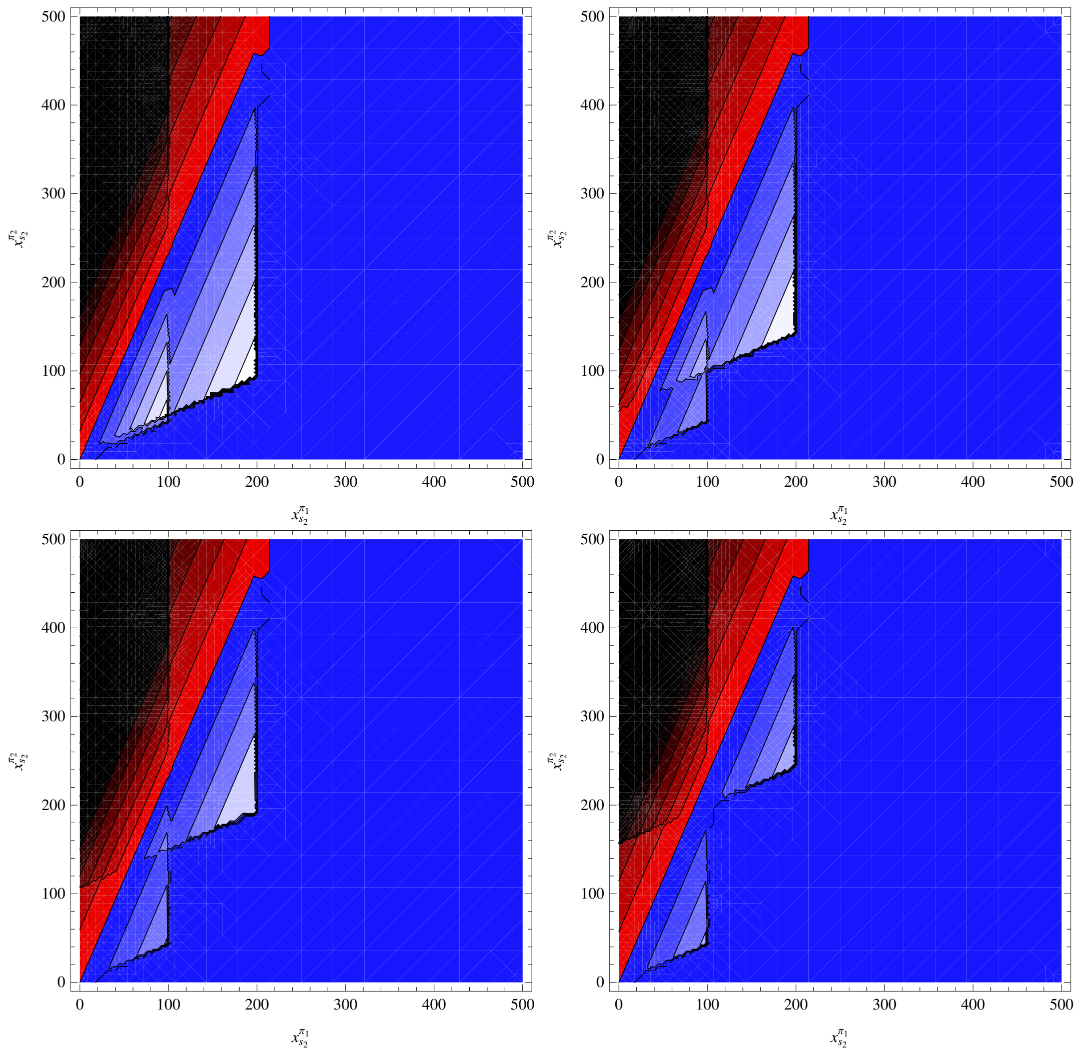

Figure 6.15: Two attribute example payoff landscape for a seller s_2, with a heterogeneous population of *bargain hunters*, Half of the buyers have budget constraints of 100.00, half of 200.00. Illustration of the effect of competition on the left peak.

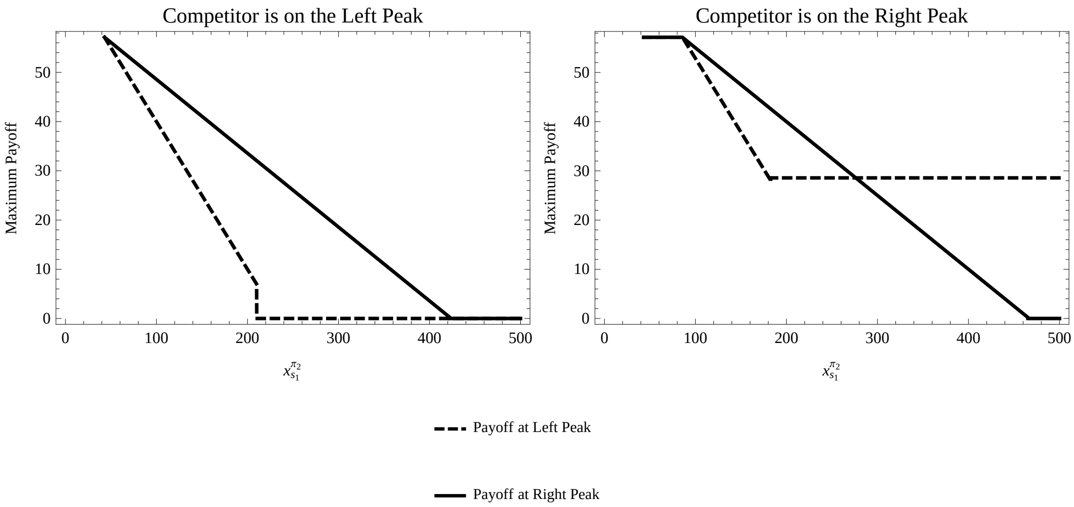

Figure 6.16: Maximum payoff available to s_2 in each market niche, as a single competitor's offer varies. Left: The competitor fixes $x_{s_1}^{\pi_1} = 100.00$ and $x_{s_1}^{\pi_2}$ varies on the x-axis. Right: The competitor fixes $x_{s_1}^{\pi_1} = 200.00$ and $x_{s_1}^{\pi_2}$ varies on the x-axis. The lines indicate the height of each peak from s_1's perspective. The buyers are *bargain hunters.*

intersect, whereas on the left plot they do not. Assuming sellers begin optimistically, offering at the top of one of the peaks near the monopolistic price, it is clear that the right peak provides more payoff regardless of which niche the competitor targets. We have seen that competition will drive their values for x^{π_2} upwards on the right peak, here towards the right on the right hand graph. This will continue until they reach the intersection point. Here, one of the sellers, say s_2, may realise that it is in its interest to switch to the left peak, where lower values for both $x_{s_2}^{\pi_1}$ and $x_{s_2}^{\pi_2}$ promise to provide a higher payoff.

Once the competitive pressure is removed from s_1, it may itself obtain a higher payoff by decreasing $x_{s_1}^{\pi_2}$ and keeping $x_{s_1}^{\pi_1}$ at 200.00. Of course, s_2 will quickly realise that it is in fact better off back on the right peak with $x_{s_2}^{\pi_1} = 200.00$ and competition on x^{π_2} will begin again.

This process is therefore cyclical, as myopic payoff-maximising agents compete against each other for the higher valued right niche, their payoffs decreasing until one of them is tempted instead by the left niche and jumps there. This relaxing of competitive pressure allows their values for the second attribute to quickly decrease, before the competition begins again.

As Kephart [9] describes, agents with different levels of strategic rationality and foresight will treat this scenario in different ways. However, most importantly it should be noted that though this kind of game theoretic reasoning may be performed here, individual sellers without sufficient knowledge of the buyers or competing sellers will not be able to do likewise, since they do not possess the details of the payoff landscape in order to find its maximum.

Therefore, what are *evolutionary market agents* to make of the landscape dynamics brought about by the *bargain hunter* behaviour? Certainly they would be expected to be able to find one or more of the peaks, though as with the single attribute case described in section 4.4, the position of the peak when they find it will be arbitrary. Similarly again to the single attribute case, there exists a large step in the payoff landscape, around two edges of both peaks. This step will cause the same issues as in the single attribute case, and leads to the same problem as described in section 4.4. The behaviour of *evolutionary market agents* in the multiple attribute model is dealt with in chapter 7.

6.5.3 Competitive Scenario: Spread Buyers

Now the above scenario, competition between two sellers for a population of buyers with heterogeneous budget constraints, is considered again, only with *spread buyers* rather than *bargain hunters*. As with the homogeneous case, for a monopolistic seller, there is no difference to the payoff landscape between *spread buyers* and *bargain hunters*. Figure 6.17 shows the example landscape presented to a monopolist.

In this particular example, both peaks are of equal height, but as shown in figure 6.11, this will not necessarily be the case. As with the previous scenarios, this analysis begins by supposing that s_1 makes an offer on the right peak of $200.00, 87.51$. This presents s_2 with the payoff landscape shown in figure 6.18.

The familiar spread buyer payoff curves appear in this landscape as in the homogeneous case, but here there is one for each market segment. It is important to note that as with the *bargain hunters* case the presence of an offer on the right peak brings down the payoff available at the left peak too. This is since a seller occupying that peak would no longer

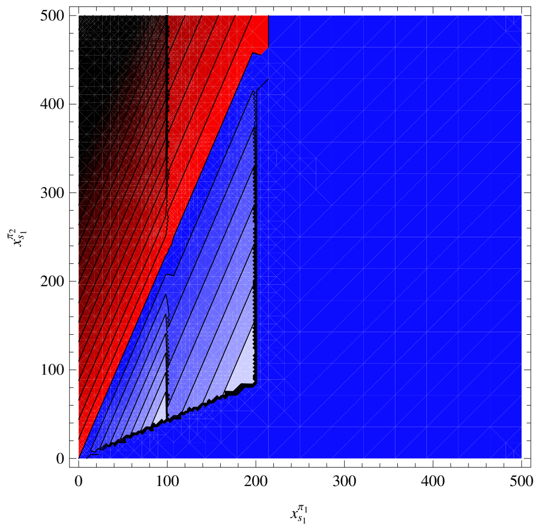

Figure 6.17: Two attribute example payoff landscape for a seller s_1, with a heterogeneous population of *spread buyers*. Half of the buyers have budget constraints of 100.00, half of 200.00. There is no competitor.

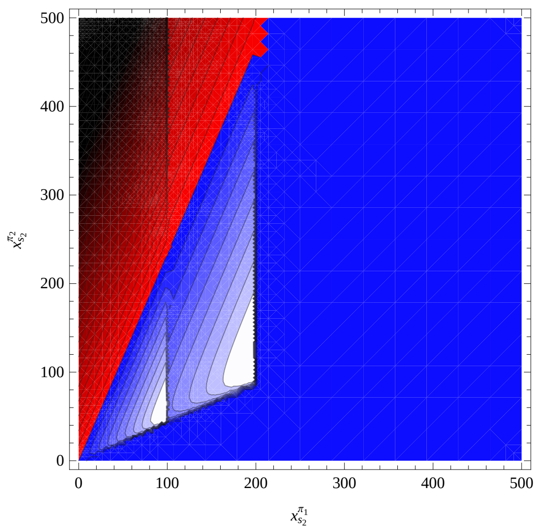

Figure 6.18: Two attribute example payoff landscape for a seller s_2, with a heterogeneous population of *spread buyers*. Half of the buyers have budget constraints of 100.00, half of 200.00. One competing seller offers $200.00, 87.51$.

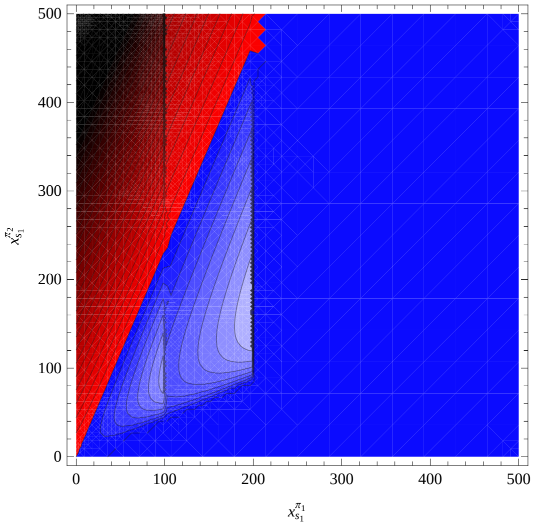

Figure 6.19: Two attribute example payoff landscape for a seller s_2, with a heterogeneous population of *spread buyers*. Half of the buyers have budget constraints of 100.00, half of 200.00. One competing seller offers $200.00, 100.00$.

be chosen by buyers with the higher budget constraint, since there is now an offer available which is more attractive to them.

There is now a global optimum from s_2's perspective, since the right peak, at $200.00, 110.14$ promises a payoff of 49.82, compared with 47.07 at the left peak. Assuming s_2 offers this, figure 6.19 shows the resulting payoff landscape for s_1.

As may be expected from the *bargain hunter* case, competition brings down the payoff available at each peak. However, unlike with *bargain hunters*, a stable equilibrium may be reached, as it is not in the sellers' interest to compete down as far as their break-even line.

Following the iterative method, in this example the two sellers are at equilibrium when one offers $100.00, 67.57$ and the other offers $200.00, 161.18$. Most importantly, with two peaks and two sellers, throughout the competitive process the sellers maximise their respective payoffs by segmenting in order to occupy different peaks. At equilibrium, the stable pair of landscapes in figure 6.20 is presented to both sellers.

In order to understand this niching effect further, figure 6.21 shows plots equivalent to those in figure 6.16 but for *spread buyers*. The plots show the height of each peak, the maximum payoff available in each niche for s_2, as its competitor s_1 increases the value of $x_{s_1}^{\pi_2}$ in its offer.

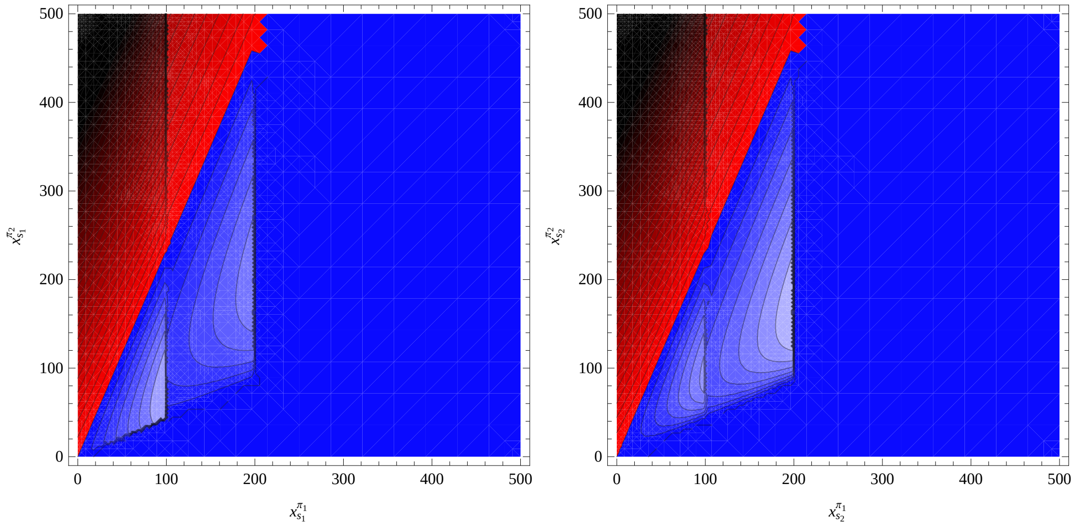

Figure 6.20: Two attribute example payoff landscapes for two sellers at equilibrium, with a heterogeneous population of *spread buyers*. Half of the buyers have budget constraints of 100.00, half of 200.00. s_1 occupies the left peak and s_2 occupies the right peak; the system is at equilibrium.

On the left, $x_{s_1}^{\pi_1}$ is fixed at 100.00, on the right at 200.00. When competition begins, as was shown in figure 6.17, the payoff at each peak was equal, however once a competing offer is made, the payoff available on the alternate peak to the competitor remains above that available on the same peak. This is the case regardless of which niche is initially targeted.

This is borne out in competition over time, as two *myoptimal* sellers update their offers to remain at their global optimum. Figure 6.22 shows the effect of competition over time on each seller's payoff.

This analysis illustrates niching behaviour between two competing sellers in a population of *spread buyers*. However, for larger number of sellers calculating the optimal decision for each seller becomes increasingly more computationally expensive, since it relies on considering the combinations of sellers providing acceptable offers to either, both or no market segments. Equally, as the number of market segments increases, a similar problem will arise. As with the single attribute model, analysing the effect of many sellers or buyer profiles on the payoff landscape quickly becomes infeasible. Therefore as the market grows, computational simulation is used instead.

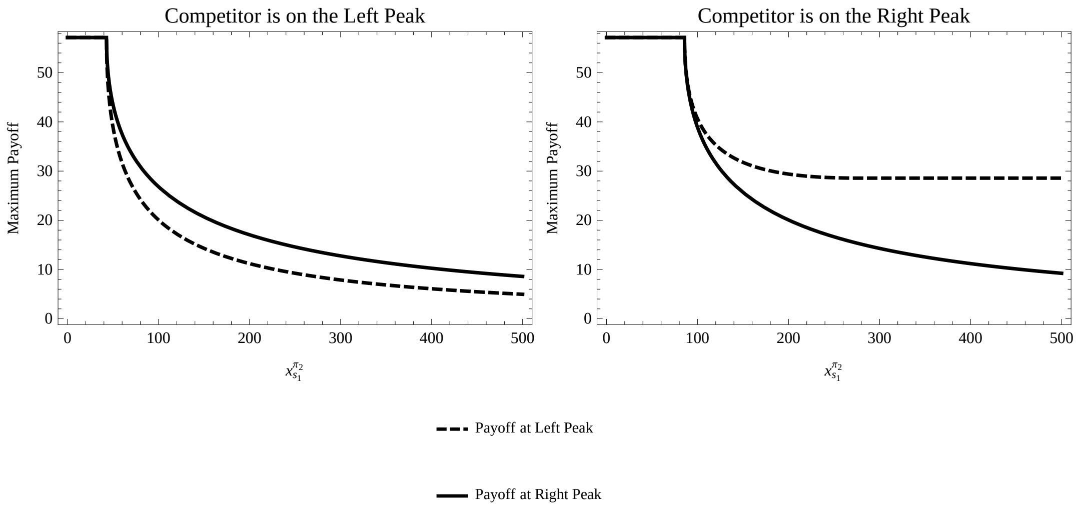

Figure 6.21: Maximum payoff available to s_2 in each market niche, as a single competitor's offer varies. Left: The competitor fixes $x_{s_1}^{\pi_1} = 100.00$ and $x_{s_1}^{\pi_2}$ varies on the x-axis. Right: The competitor fixes $x_{s_1}^{\pi_1} = 200.00$ and $x_{s_1}^{\pi_2}$ varies on the x-axis. The lines indicate the height of each peak from s_2's perspective. The buyers are *spread buyers.*

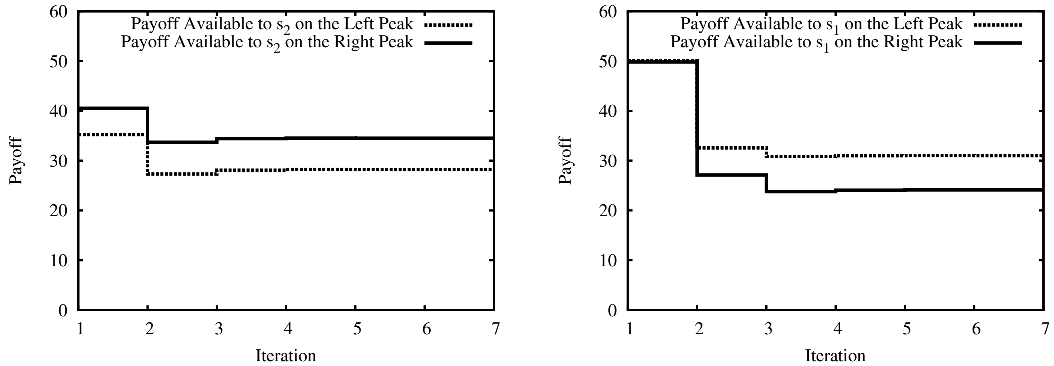

Figure 6.22: Payoff available on each peak as a result of competition between two sellers. s_2 occupies the left peak, s_1 occupies the right peak. The buyers are *spread buyers.*

6.6 Heterogeneous Buyer Preferences

Section 6.5 illustrated how market segmentation and niching may be induced by heterogeneous constraints on the acceptability of attribute values. However, market niching may also be induced by agents' heterogeneous preferences. In this thesis, the latter is termed *preference-induced niching*, which is brought about by the existence of a tradeoff between multiple desirable or undesirable attributes, between which an agent may have a preference.

At this stage it is important to highlight a difference in the treatment of price between some literature largely arising from economics and that in the field of marketing [122]. Typically in marketing price is treated as one of many attributes to be optimised as part of the *marketing mix*, in order to increase market share or maximise the profit from a sales activity. Economists however sometimes treat price as being separate from the attribute values, which describe the product being exchanged at the given price. The difference is largely unimportant here, therefore in line with more recent literature particularly on multiple attribute negotiation [123, 124], price is considered an attribute. In this case, two conflicting attributes as investigated so far, are not sufficient to allow *preference-induced niching* to occur. In order to bring about *preference-induced niching*, at least three attributes are required, once of which may of course be price.

6.6.1 Monopolistic Scenario

Considering a seller qualitatively the same as in the previous examples, its additive utility function might be extended to incorporate the third attribute, with linear attribute utility functions, such that positive values of π_1 lead to positive utility values and positive values of π_2 and π_3 lead to negative utility values. In this example its preference weights are $0.6, 0.2, 0.2$ for attributes π_1, π_2 and π_3 respectively. Again it is important to stress that this is merely an example seller used to explore the interactions between agents in the model being investigated.

Similarly, in this example the buyers' behaviour is extended without a change in their qualitative properties. A buyer's additive utility function here contains linear attribute utility functions, such that positive values of π_1 lead to negative utility values and positive values of π_2 and π_3 lead to positive utility values. Note that this is the reverse direction from the corresponding sellers, indicative of the direction of travel of the value of each attribute. In this example the buyers are divided into two equally sized groups, with preference weights $0.1, 0.1, 0.8$ and $0.1, 0.8, 0.1$ for attributes π_1, π_2 and π_3 respectively.

In both the two attribute and three attribute cases, in the absence of constraints both buyers and sellers' payoffs can be maximised by the attribute values being infinite. We saw with two attributes that given a budget constraint however, the buyers instead prefer those sellers who maximise their price within the constraint, and then provide the most attractive offer in terms of the other attribute. The same effect may be seen with three attributes.

Recall that if the buyer is given a budget constraint, such that $\overline{x}_b^{\pi_1} = 200.00$, then the offer from any seller s_i in which $x_{s_i}^{\pi_1} > 200.00$ will be rejected as unacceptable. In this case, the most attractive offer from the buyer's perspective will be that which maximises $x_{s_i}^{\pi_2}$ and $x_{s_i}^{\pi_3}$, while minimising $x_{s_i}^{\pi_1}$. However, since the buyer prefers higher values of $x_{s_i}^{\pi_3}$ more than lower values of $x_{s_i}^{\pi_1}$, due to its preference weights, a seller that chooses to increase both, in order to meet this need will find its offer preferred to one which does not. Therefore the sellers will increase in their offer the attribute value which the buyers least prefer, but which they themselves do. In this example, $x_{s_i}^{\pi_1}$ is expected to reach an equilibrium value where $x_{s_i}^{\pi_1} = \overline{x}_b^{\pi_1} = 200.00$ and for competition to ensue on the other attributes, much as was observed in section 6.4.2.

The visualisation of multiple attribute payoff landscapes when the number of attributes is greater than two is obviously more difficult than in the two attribute case in the previous sections. However, given an assumption that offers by all sellers at equilibrium will be made on the buyers' budget constraint line, where a seller s_i's offer includes $x_{s_i}^{\pi_1} = 200.00$, the visualisation of the sellers' payoff landscape may be made easier. The dimensionality of the landscape is reduced by fixing and excluding $x_{s_i}^{\pi_1}$.

Figure 6.23 therefore shows the payoff gained by a seller s_1 for selling one unit of resource as $x_{s_1}^{\pi_2}$ and $x_{s_1}^{\pi_3}$ vary, where $x_{s_1}^{\pi_1} = 200.00$. As was the case previously, lighter blue represents

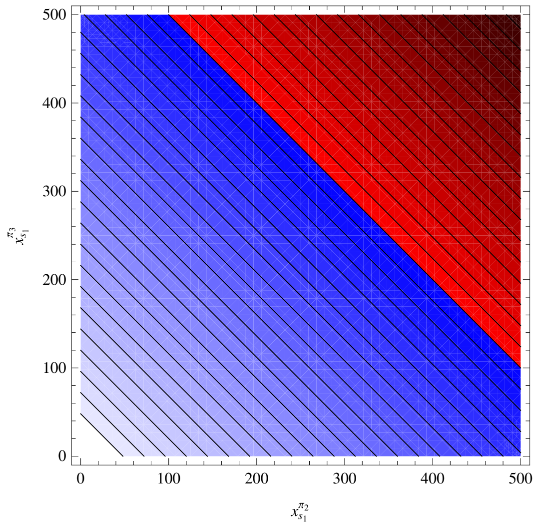

Figure 6.23: Three attribute example payoff landscape for a seller s_1 for selling one unit of resource as $x_{s_1}^{\pi_2}$ and $x_{s_1}^{\pi_3}$ vary, where $x_{s_1}^{\pi_1} = 200.00$.

areas of high payoff and red represents negative payoff, darker red being the most negative.

Of course, as with figure 6.2 this landscape assumes a unit sale. However, as described for the two attribute case in section 6.4.1, the buyers' preferences and constraints will impact upon this landscape. Figure 6.24 shows the modified payoff landscape taking account of the buyers' characteristics.

Considering a heterogeneous buyer population, where the buyer population is split into two halves with the preference weights $0.1, 0.8, 0.1$ and $0.1, 0.1, 0.8$ respectively, then these two landscapes are reduced in height by half and then added together. The landscape therefore has the appearance of those in figure 6.24 superimposed, as shown in figure 6.25.

Here there are three peaks in the payoff landscape, one specific to each buyer profile, and one which represents a best compromise between them, at the intersection of the buyers' reserve lines.

In finding this position, each buyer profile's reserve line is that which lies between the points given by the buyer's break-even value of each attribute, given a zero value of the other. More precisely, in the three dimensional attribute space, a buyer b's reserve plane will be that which is defined by the following equation:

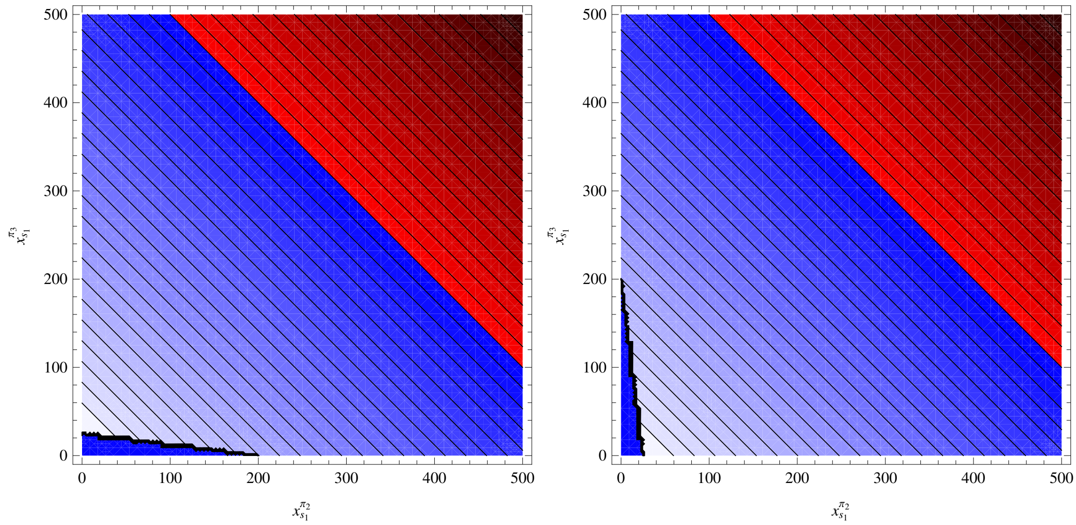

Figure 6.24: Three attribute example monopolistic payoff landscapes for a seller s_1 as $x_{s_1}^{\pi_2}$ and $x_{s_1}^{\pi_3}$ vary, where $x_{s_1}^{\pi_1} = 200.00$. Left: the buyers have weights $0.1, 0.1, 0.8$. Right: the buyers have weights $0.1, 0.8, 0.1$.

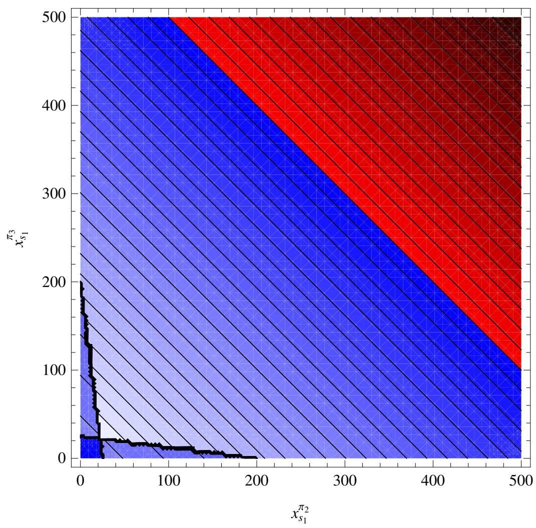

Figure 6.25: Three attribute example monopolistic payoff landscape for a seller s_1 as $x_{s_1}^{\pi_2}$ and $x_{s_1}^{\pi_3}$ vary, where $x_{s_1}^{\pi_1} = 200.00$. The buyer population is divided into two halves with preference weights $0.1, 0.1, 0.8$ and $0.1, 0.8, 0.1$ respectively.

$$k_b^{\pi_1} u_b^1(x^{\pi_1}) + k_b^{\pi_2} u_b^2(x^{\pi_2}) + k_b^{\pi_3} u_b^3(x^{\pi_3}) = 0 \tag{6.3}$$

For example, if a buyer's attribute utility functions are linear as described in section 6.2.3, weights $0.1, 0.8, 0.1$ for attributes π_1, π_2 and π_3 respectively, and $x^{\pi_1} = 200.00$, then the buyer's reserve line is therefore defined by

$$-0.02 + \frac{0.8x^{\pi_2}}{1000} + \frac{0.1x^{\pi_3}}{1000} = 0 \tag{6.4}$$

To illustrate this in terms of figure 6.25, the line in equation 6.4 intersects the horizontal axis at $x^{\pi_2} = 25.00$ and the vertical axis at $x^{\pi_3} = 200.00$.

The same method may be applied to the other buyer profile, in order to obtain the intersecting reserve line and therefore the seller's optimal offer.

6.6.2 Competitive Scenario: Bargain Hunters

In a two-seller competitive scenario, the payoff landscape will be structured similarly, but here the intersection of the two lines will be at the competitor's offer. For example, if a single competitor, s_2 were to make an offer of $200.00, 50.00, 100.00$, then the payoff landscape for s_1 would be as shown in the left plot of figure 6.26.

Dependent on the competitor's offer, though there remain three peaks in the payoff landscape, their relative heights will vary, as also shown in figure 6.26. For example, if the competitor's offer is sufficiently low, as in the left plot, then only one is a global optimum, that which is just above the intersection of the two lines. However, as shown in the right plot, as the competitor's offer moves closer to s_1's own break-even line, then the central peak loses attractiveness and the peaks are of more equal height. As the example landscapes in figure 6.27 illustrate, were s_2's offer to move closer still to s_1's break even line, one or both of the two outlying peaks may become more attractive than the central one.

Since s_2's offer is positioned exactly at the intersection of the two lines, in order to take advantage of the payoff available just above it and target the central niche, s_1 must increase

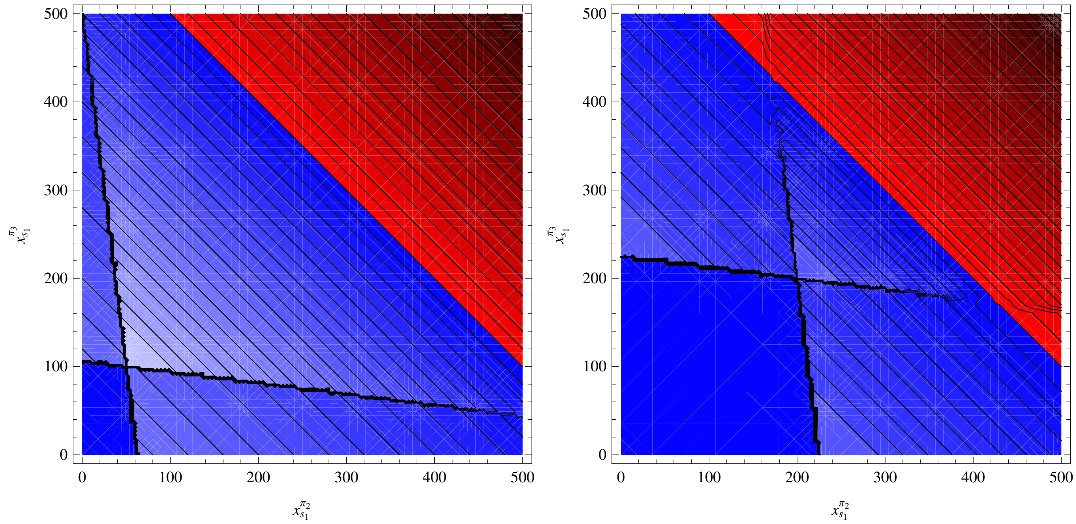

Figure 6.26: Three attribute example payoff landscapes for a seller s_1 as $x_{s_1}^{\pi_2}$ and $x_{s_1}^{\pi_3}$ vary, where $x_{s_1}^{\pi_1} = 200.00$. The buyer population consists of ***bargain hunters***, divided into two halves with preference weights $0.1, 0.1, 0.8$ and $0.1, 0.8, 0.1$ respectively. The lines intersect at the competitor's offer. Left: s_2 offers $200.00, 50.00, 100.00$. Right: s_2 offers $200.00, 200.00, 200.00$.

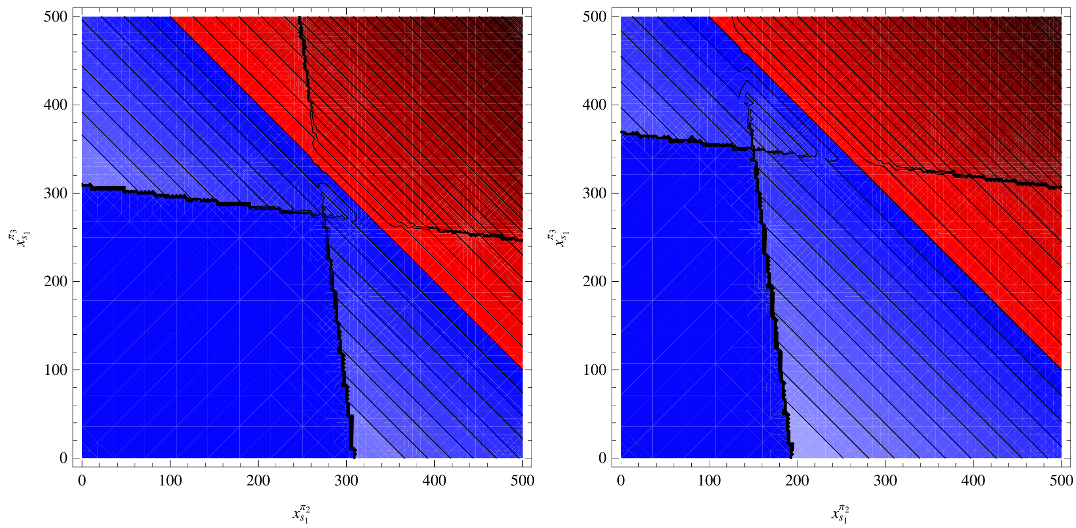

Figure 6.27: Three attribute example payoff landscapes for a seller s_1 as $x_{s_1}^{\pi_2}$ and $x_{s_1}^{\pi_3}$ vary, where $x_{s_1}^{\pi_1} = 200.00$. The buyer population consists of ***bargain hunters***, divided into two halves with preference weights $0.1, 0.1, 0.8$ and $0.1, 0.8, 0.1$ respectively. The lines intersect at the competitor's offer. Left: s_2 offers $200.00, 275.00, 275.00$. Right: s_2 offers $200.00, 150.00, 350.00$.

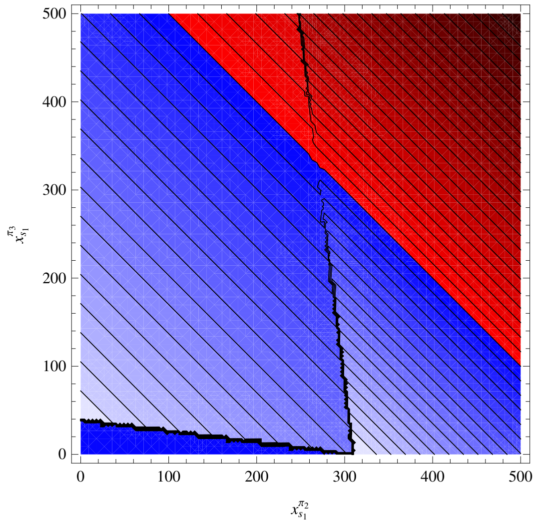

Figure 6.28: Three attribute example payoff landscape for a seller s_1 as $x_{s_1}^{\pi_2}$ and $x_{s_1}^{\pi_3}$ vary, where $x_{s_1}^{\pi_1} = 200.00$. The buyer population consists of ***bargain hunters***, divided into two halves with preference weights $0.1, 0.1, 0.8$ and $0.1, 0.8, 0.1$ respectively. A competitor, s_2 offers $200.00, 310.00, 0.00$.

either $x_{s_1}^{\pi_2}$ or $x_{s_1}^{\pi_3}$ or both by a small amount. Of course, from s_2's perspective, the intersection of the lines is then shifted to the position of s_1's offer, forcing it to do the same in order to remain in this niche. The intersection moves further away from the origin, and as it does so its attractiveness decreases with respect to the two outlying niches. This competition will continue until either the payoff available in the central niche is zero, or else low enough to trigger a seller to move to one of the outlying peaks.

Once a seller chooses instead to target one of the outlying niches, its competitor is immediately encouraged to do likewise. Consider that s_2 makes an offer at $200.00, 310.00, 0.00$, targeting the lower niche. The payoff landscape for s_1 is as shown in figure 6.28.

This move shifts the intersection of the lines back away from the sellers' break-even line. However, as before, as competition between the sellers ensues on this peak, the intersection point shifts away from the origin and payoff decreases. Figure 6.29 illustrates this.

As is now clear, the reduction of payoff available on the peak being competed for leads to an increasing difference in payoff between the two niches. Eventually a seller will once again switch to the other peak, and the process will begin again. Hence, regardless of which peak is initially chosen by the sellers, the cyclical competition wars will continue.

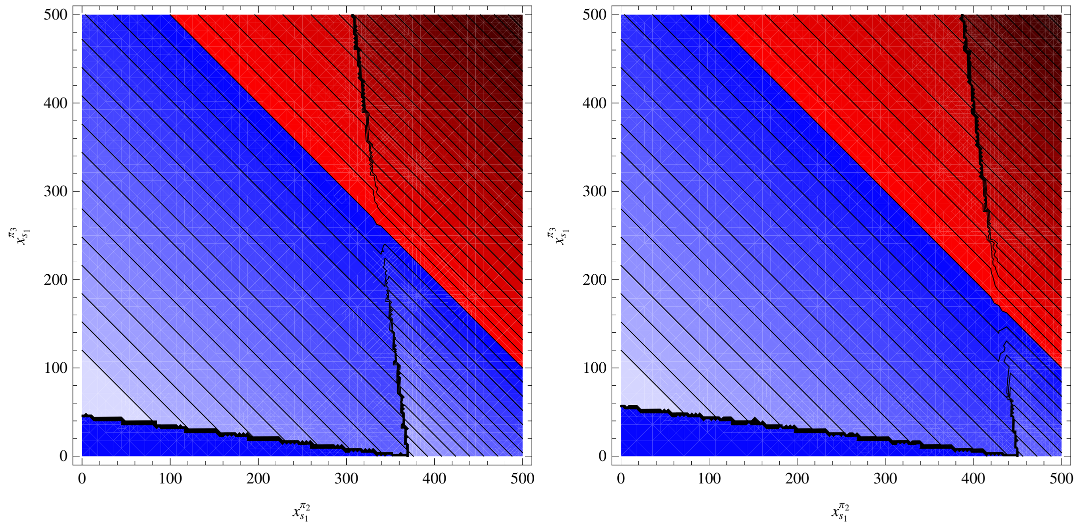

Figure 6.29: Three attribute example payoff landscapes for a seller s_1 as $x_{s_1}^{\pi_2}$ and $x_{s_1}^{\pi_3}$ vary, where $x_{s_1}^{\pi_1} = 200.00$. The buyer population consists of ***bargain hunters***, divided into two halves with preference weights $0.1, 0.1, 0.8$ and $0.1, 0.8, 0.1$ respectively. Competition ensues on the lower peak as s_2 offers $200.00, 370.00, 0.00$ (left) and later $200.00, 450.00, 0.00$ (right).

As with section 6.5, when buyers act as ***bargain hunters***, the presence of more than one peak in the sellers' payoff landscape ensures that their myopic behaviour leads to an unending cycle of competitive undercutting of offers. This again is an example of the limit-cycle phenomenon observed by Kephart et al. [9].

6.6.3 Competitive Scenario: Spread Buyers

Similarly to sections 6.4 and 6.5, the multiple attribute model is now considered where buyers with heterogeneous preferences act as ***spread buyers*** rather than ***bargain hunters***.

Clearly as before, the difference in the buyers' decision function will not affect the outcome in the monopolistic case. However, once a competitor is present, the ***spread buyer*** behaviour will manifest itself through the sellers' payoff landscapes. Accordingly, figure 6.30 shows the payoff landscape for a seller s_1, when responding to a competitor s_2's offer of $200.00, 200.00, 200.00$. The buyers are ***spread buyers***, and the landscape is shown for the cases when the buyers' weights are $0.1, 0.1, 0.8$ and $0.1, 0.8, 0.8$. As with ***bargain hunters*** it is assumed that sellers will offer the equilibrium value for x^{π_1} of 200.00.

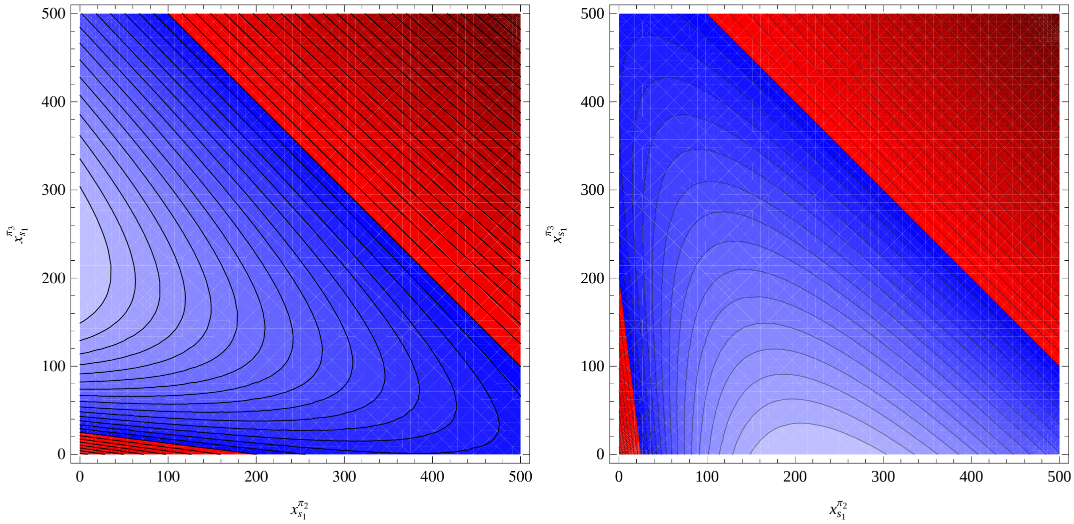

Figure 6.30: Three attribute example payoff landscapes for a seller s_1 as $x_{s_1}^{\pi_2}$ and $x_{s_1}^{\pi_3}$ vary, where $x_{s_1}^{\pi_1} = 200.00$. One competitor s_2 offers $200.00, 200.00, 200.00$. In each case the buyers are a homogeneous population of *spread buyers*. Left: their preference weights are $0.1, 0.1, 0.8$. Right: their preference weights are $0.1, 0.8, 0.1$.

As with the *bargain hunter* case, splitting the buyers into two equally sized profiles, with preference weights $0.1, 0.1, 0.8$ and $0.1, 0.8, 0.1$ respectively, the payoff landscapes may be thought of as being superimposed. Figure 6.31 shows the resulting payoff landscape for s_1.

Perhaps the most notable result of the buyers switching to *spread buyer* behaviour is that there are no longer multiple peaks or niches in the sellers' payoff landscape. Instead, a large central peak appears, at which the seller can maximise its payoff.

Indeed in this example, a *myoptimal* seller maximising its payoff would be expected to offer at the optimal market position of $200.00, 129.03, 129.03$. Assuming s_2 did so, s_1 is then faced with the landscape in figure 6.32.

As is now well rehearsed for *spread buyers*, an iterative process may be followed in order to determine the equilibrium point in the attribute space. In this example, equilibrium is reached when both sellers offer $200.0, 114.81, 114.81$. At this point, neither seller can increase its payoff by moving its offer from this point, as illustrated by figure 6.33.

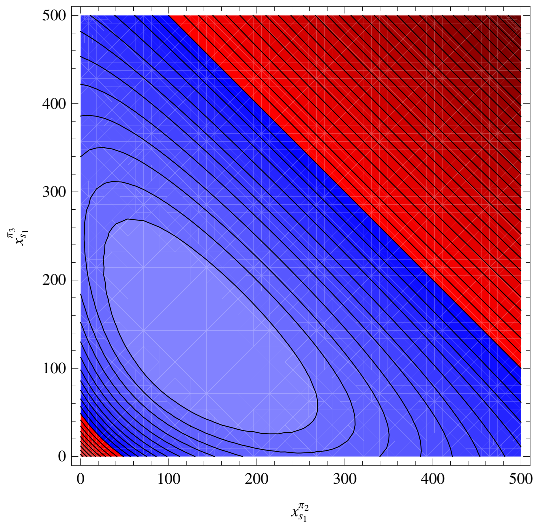

Figure 6.31: Three attribute example payoff landscape for a seller s_1 as $x_{s_1}^{\pi_2}$ and $x_{s_1}^{\pi_3}$ vary, where $x_{s_1}^{\pi_1} = 200.00$. The buyer population consists of *spread buyers*, divided into two halves with preference weights $0.1, 0.1, 0.8$ and $0.1, 0.8, 0.1$ respectively. One competitor s_2 offers $200.00, 200.00, 200.00$.

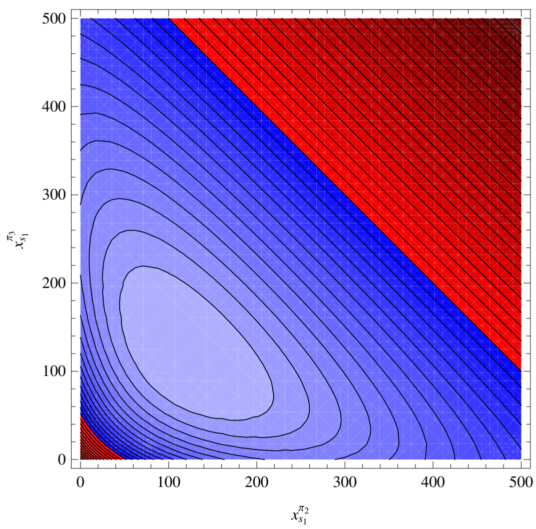

Figure 6.32: Three attribute example payoff landscape for a seller s_1 as $x_{s_1}^{\pi_2}$ and $x_{s_1}^{\pi_3}$ vary, where $x_{s_1}^{\pi_1} = 200.00$. The buyer population consists of *spread buyers*, divided into two halves with preference weights $0.1, 0.1, 0.8$ and $0.1, 0.8, 0.1$ respectively. One competitor s_2 offers $200.00, 129.03, 129.03$.

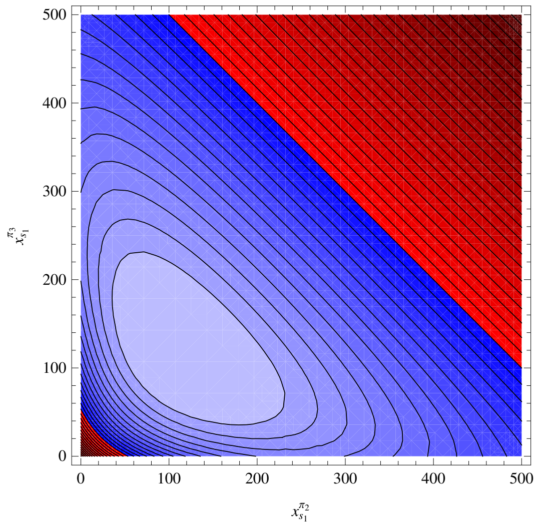

Figure 6.33: Three attribute example payoff landscape for a seller s_1 as $x_{s_1}^{\pi_2}$ and $x_{s_1}^{\pi_3}$ vary, where $x_{s_1}^{\pi_1} = 200.00$. The buyer population consists of ***spread buyers***, divided into two halves with preference weights $0.1, 0.1, 0.8$ and $0.1, 0.8, 0.1$ respectively. One competitor s_2 offers $200.00, 114.81, 114.81$. The system is at equilibrium.

6.7 Conclusions and Discussion

In this chapter, the single attribute posted offer market model as introduced in chapter 3 was extended, in order that resources may be described over additional quality attributes. Resource users may have different preferences and constraints over the attributes, and as such form niches in the market, into which providers may segment.

In scenarios where seller payoff landscapes have multiple peaks, brought about by either heterogeneous buyer constraints or by heterogeneous buyer preferences, ***myoptimal*** seller behaviour leads to limit-cycle wars, as observed in the qualitatively similar model studied by Kephart et al. [9]. They argue that the combination of three factors cause these unstable dynamics, firstly multiple peaks in the payoff landscape as discussed, secondly the ability of agents to possess sufficient knowledge of the landscape in order to identify where the global optimum lies as well as move immediately to it, and thirdly the inability of myopic agents to anticipate the response of their competitors. Given the example scenarios investigated in this chapter, a fourth factor appears to be the presence of hyperrational *bargain hunter* style behaviour on the part of buyers. Buyer behaviours which give rise to smoother seller payoff functions, such as the *spread buyers* investigated here, can mitigate against the limit-cycle war effect and instead provide for a stable equilibrium.

Therefore, this chapter makes the following contributions:

- An extension of the posted offer market-based approach to resource allocation, to consider resources which may be described over multiple quality attributes in addition to or in place of price.

- A potential extension to the buying and selling agents' decision models, grounded in multi-attribute utility theory, which enables them to consider the additional attributes present in the extended model.

- A description of two potential causes of market segmentation, namely *constraint-induced niching*, due to buyers' heterogeneous constraints and *preference-induced niching*, due to their heterogeneous preferences over attributes.

- Analysis of the dynamics of the system in a range of representative example scenarios, where buyers are either *bargain hunters* or *spread buyers.* Limit-cycle or price wars are observed where buyers act as *bargain hunters*, however their behaviour as *spread buyers* mitigates this and instead leads to a stable, predictable equilibrium.

As with the differences between chapters 3 and 4, the idealised behaviour described by Kephart et al. [9] and assumed in the analysis in this chapter, cannot be expected to be wholly borne out in practical scenarios. The following chapter, chapter 7 will therefore investigate the extended multiple attribute model where sellers have access to only private information, by means of evolutionary simulation.

CHAPTER 7

A DIVERSITY DILEMMA IN EVOLUTIONARY MARKETS

The man who follows the crowd will usually get no further than the crowd. The man who walks alone is likely to find himself in places no-one has ever been.

Alan Ashley-Pitt

This chapter investigates the extended multiple attribute market model introduced in chapter 6, in simulation, through the use of *evolutionary market agents* introduced in chapter 4. Through evolutionary simulations, it is shown that resource providing nodes using *evolutionary market agents* are able to position themselves in the niches predicted in chapter 6, thereby providing levels of quality attributes appropriate to the users. However, it is further shown how algorithmic artifacts may distort this to varying degrees, due to a dilemma faced by evolutionary sellers between a maintaining high population diversity to perpetuate search and a low diversity to exploit a market position and hence gain a higher payoff. This dilemma is illustrated using a parameterised hypermutation operator.

Of course, the dilemma between maintaining high diversity for exploration and low diversity for exploitation is not new. Indeed, in a discrete choice form it is often characterised in terms of a family of *bandit* problems, which quickly become very difficult to solve analytically, even for relatively small cases [125]. The problem is also well known in the evolutionary computation community [126] and according to Holland a classic problem affecting all adaptive and learning systems [127].

However, what is investigated here is different from those evolutionary cases previously studied, due to the online nature of the fitness evaluation. Particularly, in classical evolutionary search and optimisation, the goal is that one individual within the population finds the optimum. In the case studied here, performing a fitness evaluation on a given individual in the population counts towards the total fitness which is obtained by the algorithm over time. The goal in optimising the tradeoff presented by the diversity dilemma exposed in this chapter is to maximise total payoff over time, calculated across all fitness evaluations.

The remainder of the chapter is organised as follows. Section 7.1 demonstrates through the use of simulation results, that *evolutionary market agents*, as introduced in chapter 4 and without modification are able to position sellers in market niches, by considering a conjoint utility value for each offer as that offer's fitness value in the evolutionary algorithm. In section 7.2 the number of sellers finding each niche is examined, and it is demonstrated that the standard *evolutionary market agent* algorithm can lead to uneven segment sizes. Section 7.3

describes the population diversity dilemma which explains why the uneven segments occur and sellers are unable to obtain the optimal payoff. The dilemma is explored in section 7.4 by means of a parameterised hypermutation operator, which frames the dilemma as a tradeoff in the algorithm's design parameters. Section 7.5 examines what behaviour might be expected of self-interested sellers in the context of the diversity dilemma. Finally, section 7.6 concludes the chapter with a discussion.

7.1 Demonstrating Market Segmentation

We have seen so far how the model introduced and investigated in earlier chapters may be extended to take account of multiple attributes, and we have further seen the effect of multiple attributes on market dynamics with *myoptimal* sellers. However, as discussed in chapter 4, the use of *myoptimal* sellers in this analysis makes the assumption that sellers have access to global information about their payoff landscape, which in turn requires up to date knowledge of all other agents in the system, including the buyers' utility functions. This is considered private information, which is not generally available. The *evolutionary market agent* algorithm however is a heuristic approach, which makes use only of local, private information.

7.1.1 Experimental Setup

In this section, as with chapter 5, the focus is on *spread buyer* behaviour. Firstly, this is since it was shown in chapter 4 that the combination of *evolutionary market agents* and *bargain hunters* leads to instability, and secondly that in the multiple peak landscapes arising from the multiple attribute model, *bargain hunters* lead to a cyclical competitive dynamic, known as limit-cycle wars between sellers. By contrast, *spread buyers* were shown to lead to stable and predictable outcomes, due to the smooth nature of their payoff functions.

In the context of the *evolutionary market agent* algorithm, the conjoint utility measure provided by the multi-attribute utility function is therefore seen as the payoff associated with a certain offer, and hence used as a fitness measure for each agent's evolutionary algorithm.

The initial experimental setup consisted of 20 seller agents and 1000 buyer agents trading a resource π with two attributes, π_1 and π_2, notionally price and quality respectively. Both buyer and seller populations had homogeneous linear attribute utility functions, such that attribute π_1 was desirable to the sellers but undesirable to the buyers, while attribute π_2 was undesirable to the sellers but desirable to the buyers. In terms of the additive utility model's scaling constants, sellers were given the weights 0.7, 0.3, while buyers were given the weights 0.3, 0.7. The exact values of these weights, and their symmetry, proved not to be important, however they were found to provide a good range within the attribute space for feasible offers.

7.1.2 Constraint-induced Niching

Constraints, as introduced in section 6.5, were used to generate market segmentation or niching. Each buyer had a hard budget constraint, in terms of attribute π_1, notionally price, such that any offer for which x^{π_1} was greater than this constraint is rejected as unacceptable by that buyer. The value of each buyer's budget constraint was drawn with equal probability from one of two Gaussian distributions, with a relatively low standard deviation when compared with the offer space as a whole. The means of these distributions were 100.00 and 200.00. This is similar to the scenario investigated in section 6.5, when heterogeneous buyer constraints are first studied.

Figure 7.1 shows the position of each of the 20 sellers' offers in the attribute space, at a given point in time. Recall that each point, each offer in the attribute space is the result of an independent evolutionary algorithm, and that no information is shared between algorithms. The vertical lines in each plot represent the buyers' budget constraints in terms of attribute x^{π_1}, which varies horizontally. The upper and lower diagonal lines represent the break even lines for the sellers and buyers respectively. These features are derived from their utility functions and preference weights as described in chapter 6, and may be compared with the features on, for example, the payoff landscape in figure 6.17.

As can be clearly seen, from an initial random distribution, sellers' offers quickly found the feasible region between the two diagonal lines. The offers then segmented, converged to two niches, one for each subset of the buyer population. The system subsequently remains stable, other than for ongoing small mutations about the centre of each niche. In summary, the seller agents self-organise to the predicted equilibrium values for their offers' attributes. The final plot in figure 7.1 may be compared with the peaks in the equilibrium payoff landscapes shown in figure 6.20.

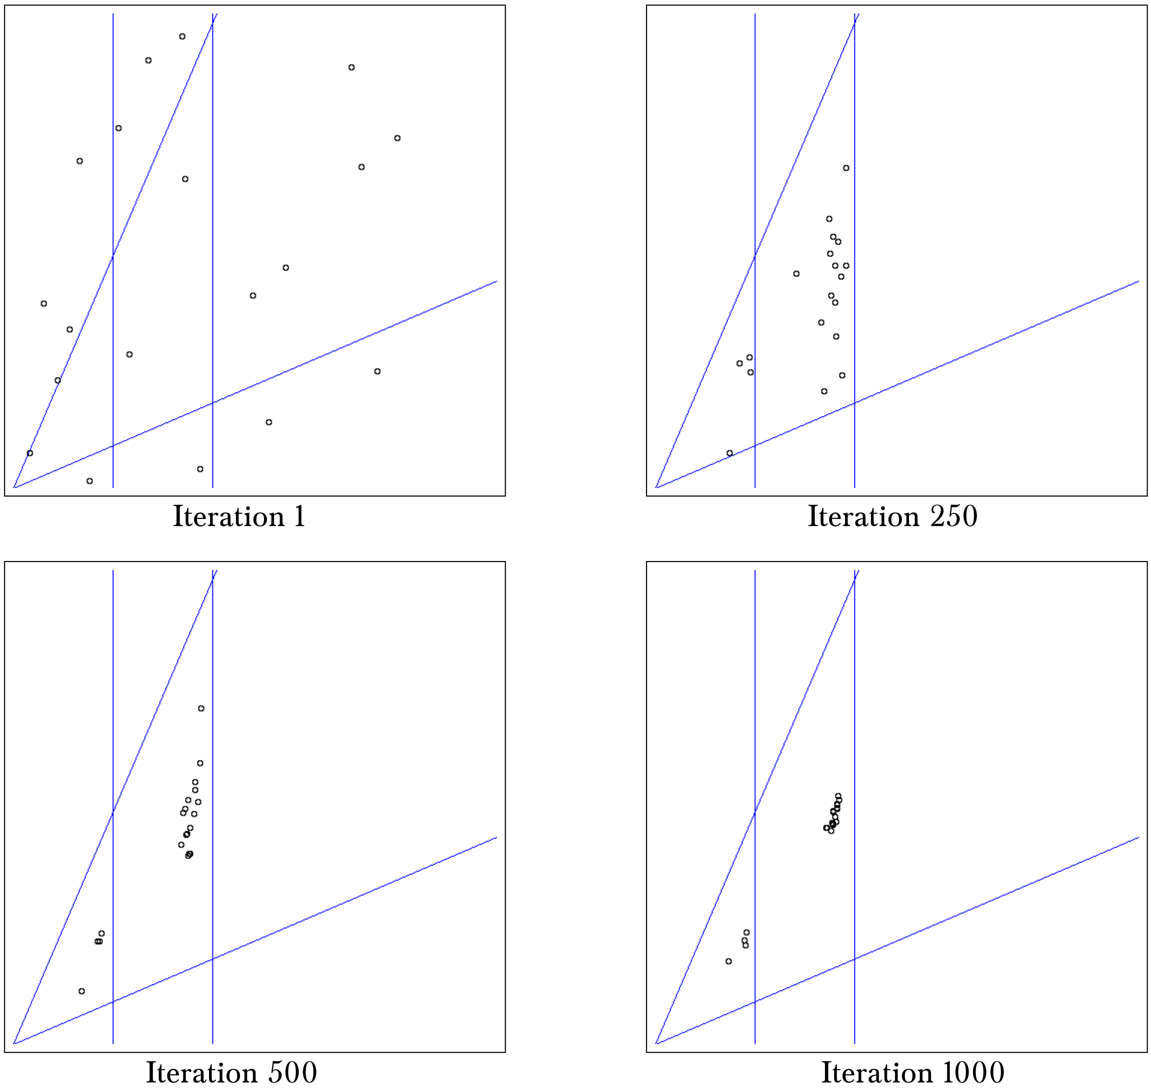

Figure 7.1: Evolution of the offers in the attribute space through time for a typical simulation run. Each circle represents an individual seller's offer. The space between the diagonal lines is the feasible region in which offers are acceptable to both the buyers and sellers. The two vertical lines indicate the means of the Gaussian distributions from which the buyers' budget constraints are drawn. Evolutionary market agents lead the sellers to find the niches created by the buyers' heterogeneous budget constraints.

7.1.3 Preference-induced Niching

Chapter 6 also described how market segmentation may be brought about by heterogeneous preferences on the part of buyers, rather than directly through budget constraints. In this next series of simulations, the scenario described in section 6.6 is reproduced, where a third attribute π_3 is added. Here, sellers make use of the additive utility model, with linear attribute utility functions, such that positive values of π_1 lead to positive utility values and positive values of π_2 and π_3 lead to negative utility values. In this example the sellers' preference weights are $0.6, 0.2, 0.2$ for attributes π_1, π_2 and π_3 respectively. Buyers also make use of the additive utility model with linear attribute utility functions, such that positive values of π_1 lead to negative utility values and positive values of π_2 and π_3 lead to positive utility values.

In the first baseline simulation, buyers are homogeneous with preference weights $0.1, 0.1, 0.8$ for attributes π_1, π_2 and π_3 respectively. Figure 7.2 shows the evolution of the sellers' offer for a typical simulation run with this homogeneous population of buyers.

The plots in this figure can be compared with with the payoff landscape shown in figure 6.30, where the optimal offer is for the sellers to minimise the buyers' least preferred attribute, x^{π_2} and instead compete to provide x^{π_3}. This is indeed what is observed in simulation, as shown in figure 7.2, and a stable equilibrium is reached.

Interestingly, as section 6.6 illustrated, when the buyer population has heterogeneous preferences, *spread buyer* behaviour rather than bringing about market segmentation as with *bargain hunters*, instead led to an equilibrium payoff landscape where the optimal offer was a compromise which targeted neither of the two buyer profiles fully. As with the analysis in section 6.6, in the next simulation the buyer population is divided into two equally sized groups, with preference weights $0.1, 0.1, 0.8$ and $0.1, 0.8, 0.1$ for attributes π_1, π_2 and π_3 respectively.

Figure 7.3 shows the evolution over time of the sellers' offers for a typical run of this simulation. As with the homogeneous case in figure 7.2, the *evolutionary market agents* lead the sellers to reach equilibrium as predicted in the landscape analysis presented in section 6.6. Of note also in this particular example is that the sellers' offers do not converge tightly to the peak in the payoff landscape. Due to the lack of a sharp peak, as seen in the other

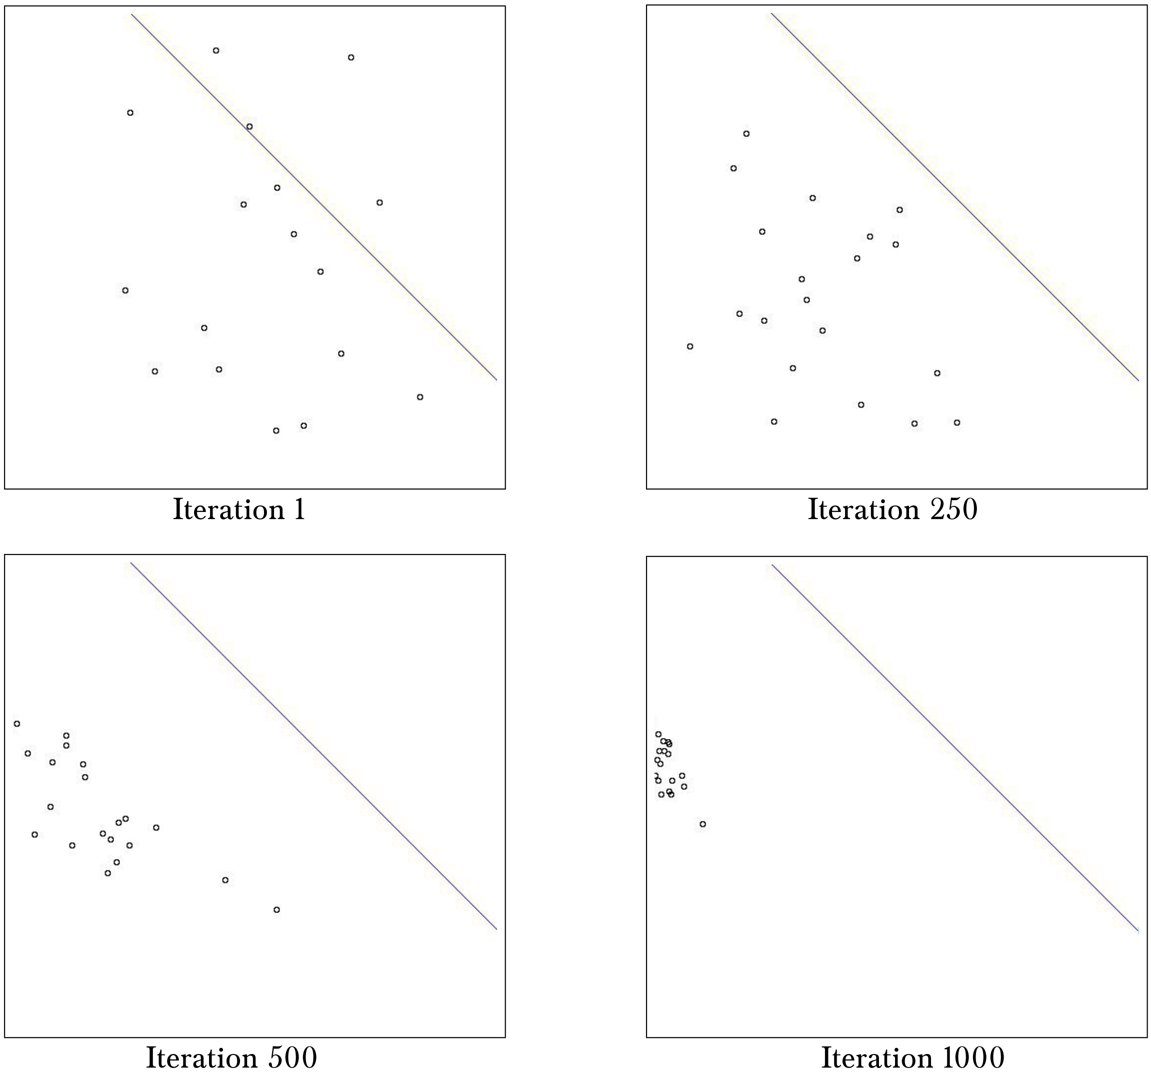

Figure 7.2: Evolution of the offers in the attribute space through time for a typical simulation run with three attributes. Attribute x^{π_2} is on the horizontal axis, x^{π_3} on the vertical. Each circle represents an individual seller's offer. The blue line represents the sellers' break-even line when $x^{\pi_1} = 200.0$. The buyers are homogeneous with weights $0.1, 0.1, 0.8$. Evolutionary market agents lead the sellers to find the equilibrium offer, where provision is focused on attribute x^{π_3} at the expense of x^{π_2}, according to the buyers' weights.

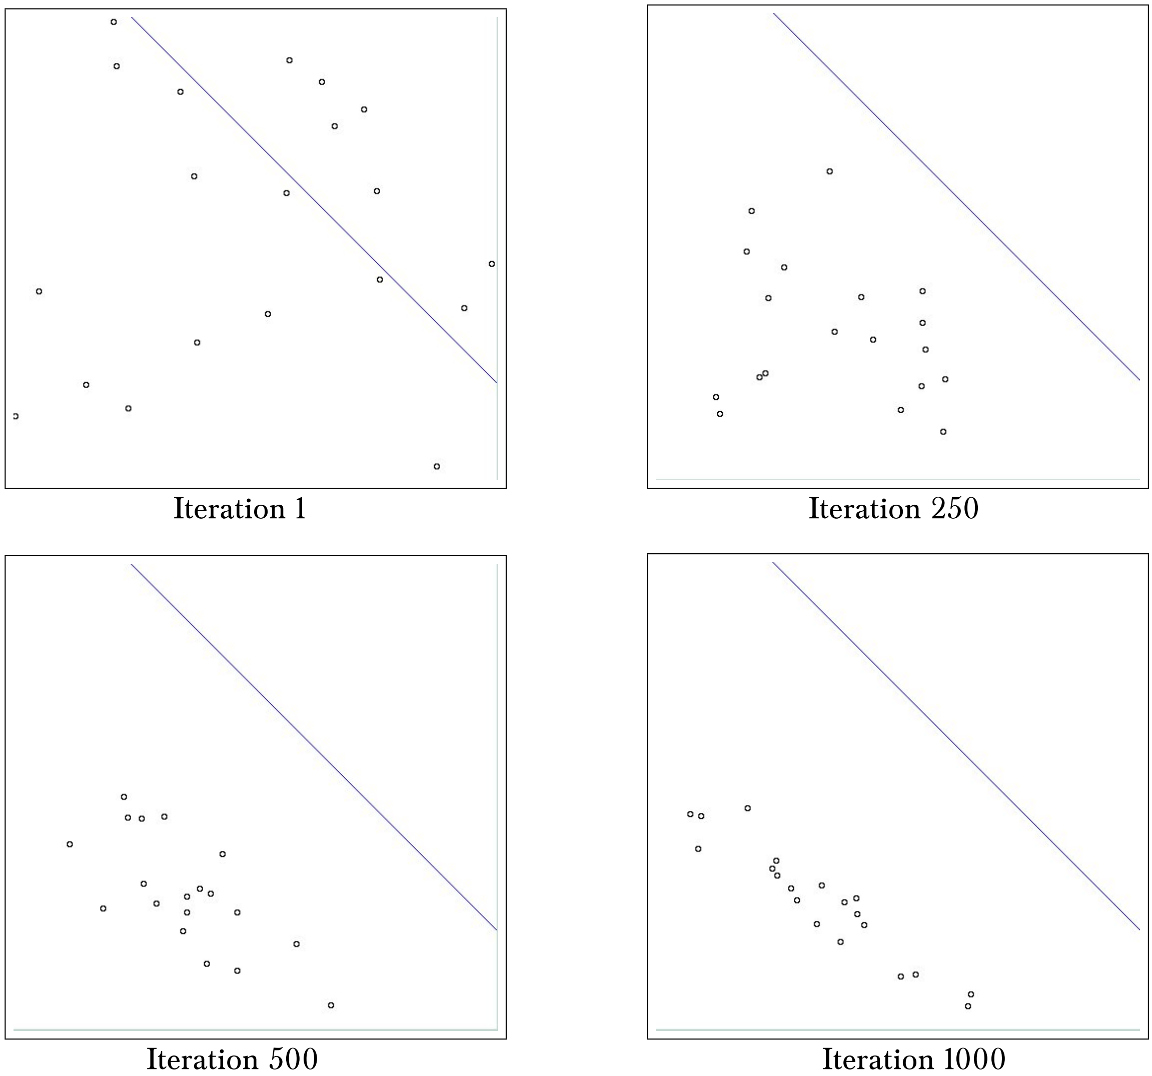

Figure 7.3: Evolution of the offers in the attribute space through time for a typical simulation run with three attributes. Attribute x^{π_2} is on the horizontal axis, x^{π_3} on the vertical. Each circle represents an individual seller's offer. The blue line represents the sellers' break-even line when $x^{\pi_1} = 200.0$. The buyers are divided into two equal profiles, with weights $0.1, 0.1, 0.8$ and $0.1, 0.8, 0.1$ respectively. Evolutionary market agents lead the sellers to reach equilibrium, where offers represent a compromise between both buyer profiles.

landscapes, they instead remain relatively spread out across the plateau which exists in this particular example.

Finally, the assumption that sellers would always offer a value of x^{π_1} at the buyers' budget constraint, which underlies the validity of these figures, should be checked. In all the simulations in this section, this was found indeed to be the case. Figure 7.4 shows the evolution of sellers' offers for x^{π_1} across the 30 independent runs of the final simulation. The budget constraint, which serves as the equilibrium value for x^{π_1} was 200.00.

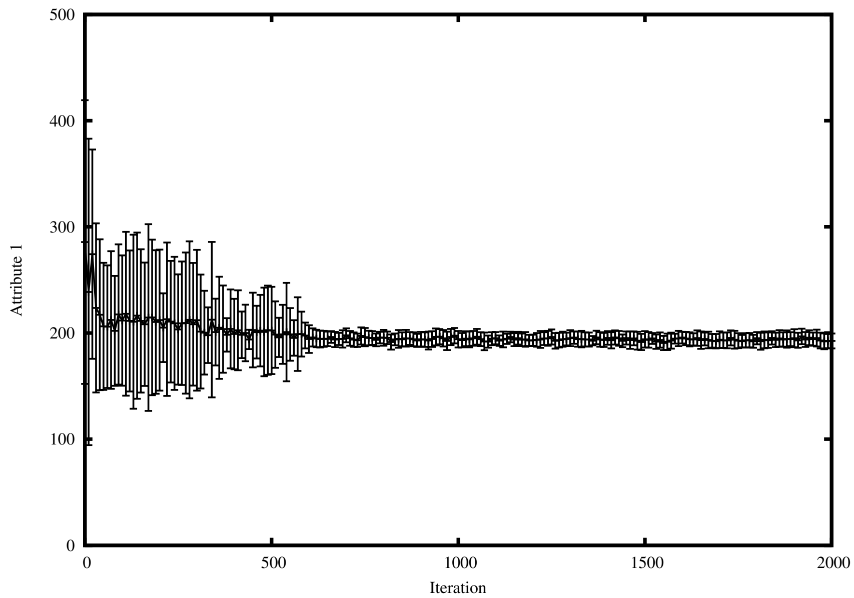

Figure 7.4: Evolution of mean attribute x^{π_1} values through time in the simulation depicted in figure 7.3. Mean and standard deviation shown over 30 independent runs. The sellers' offers find the buyers' budget constraint and remain there.

7.2 Quantifying Market Segmentation

Section 7.1 showed that sellers' use of *evolutionary market agents* allows them to find and position their offers market niches that are otherwise unknown to them. However, it is also clear from the results in figure 7.1 that the sellers are not evenly split between the niches. In this and the following sections, the behaviour of sellers in the scenario described in section 7.1.2 is investigated further. Figure 7.5 shows for that experiment the number of sellers positioned

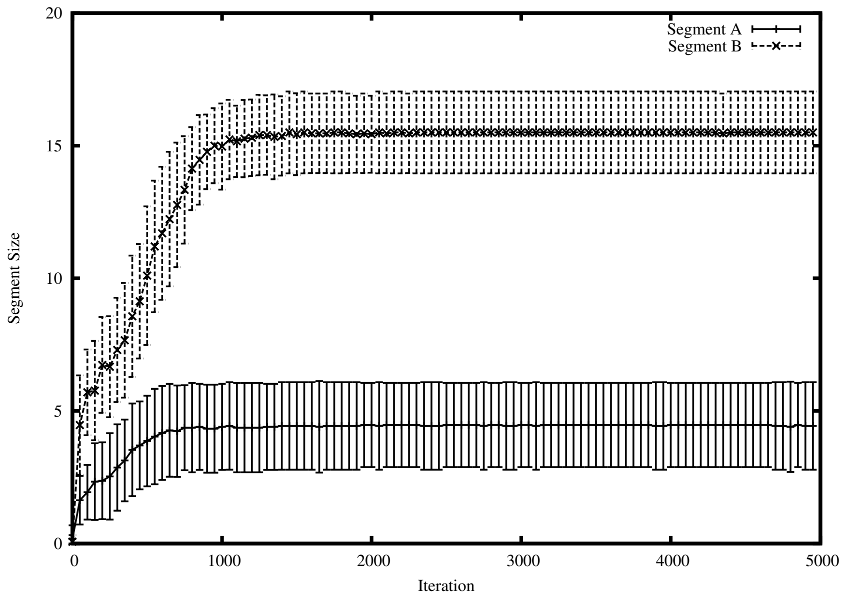

Figure 7.5: Number of sellers in each market segment. Segment A, that at the left niche, repeatedly contains fewer sellers than segment B, at the right niche.

within a radius of 30 from the centre of each niche for this experiment, as the simulation progresses. Mean and standard deviation are shown for 30 independent runs.

That this is not optimal seller behaviour can be verified by looking at the payoffs the sellers obtain in each niche. Figure 7.6 shows the mean payoffs, along with their standard deviation, for the sellers in each segment, over time. The sellers in the left segment, call it segment A are generally obtaining a higher payoff than those in the right segment, segment B. However, their payoff is highly unpredictable, indicated by the very high standard deviation across the independent runs. Conversely, the sellers in segment B are obtaining a lower payoff, though are highly stable over time.

The question then arises of why the sellers in segment B do not move their offers to the niche occupied by segment A, given the increased likelihood of obtaining a higher payoff for doing so. Indeed, beyond iteration 1000, there is relatively little change in either segment size or payoff. This is due to the lack of diversity in the sellers' populations of offers after this stage in the simulation. Since the populations are highly converged, making a jump to the more attractive niche is very unlikely.

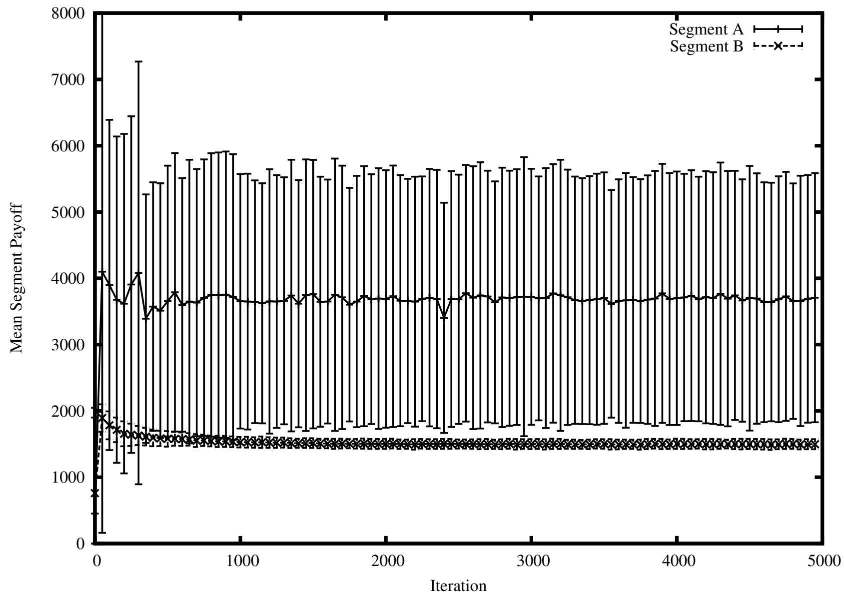

Figure 7.6: Mean and standard deviation seller payoffs for each segment. The payoff for sellers in segment A, the lower valued niche, is greater though highly unpredictable. The payoff for segment B is lower though more stable.

7.3 The Diversity Dilemma

Of vital importance when considering the behaviour of evolutionary sellers in this simulation is that since the seller has no model of the market, in order to obtain a fitness value for an offer in its population, it must test this offer in the live market. This is due to information being private and therefore a seller has no knowledge of the buyers or behaviour of its competitors, other than that which may be inferred through the payoff returned from its actions. Performing a fitness evaluation on a given offer in the population therefore counts towards the total payoff which is obtained by the seller over time.

Since every fitness evaluation counts, once a seller can be sure it has found the best position for its offer, it can maximise its payoff over time by always taking that position. In terms of the evolutionary algorithm, this means the seller's population of offers converging to that point, losing all diversity.

However, a seller can never be sure it has found this position, since even if it is at a local optimum, a better position may exist in the attribute space. This is what we observe occurring in the simulation results described in section 7.2. In order for a seller to find a position in the attribute space which would lead to a higher payoff, or indeed to gain confidence that there does not exist such a position, the seller must perform a search. In the context of the evolutionary algorithm, this search requires population diversity.

Since the evolutionary market agent algorithm has been hitherto tested primarily in single attribute environments with single equilibria, it is no surprise that the algorithm favours an initial search period with high population diversity, in which the equilibrium is found, followed by a loss of diversity in the population, as the discovered position may be exploited to obtain maximum payoff. In scenarios such as that investigated in this chapter however, where from the perspective of the algorithm more than one peak exists, a different approach to the diversity dilemma may be needed.

7.4 Exploring the Diversity Dilemma

The focus in this section is to explore the diversity dilemma described in section 7.3, and attempt to mitigate its effect on the overall payoff which a seller achieves. As discussed at the beginning of this chapter, the problem of population diversity management in evolutionary computation is not new, and indeed increased population diversity is a characteristic which is often found to be desirable in evolutionary algorithms applied to dynamic optimisation problems [78]. Methods of diversity preservation include immigrant schemes, random restarts and hypermutation [128]. This thesis does not claim that any particular approach is any better suited to the diversity dilemma investigated here than any other, however in this section an example of a diversity preservation technique is used in order to characterise the diversity dilemma. Specifically, hypermutation is adopted as it has been shown to preserve diversity without discarding as much information as alternative schemes such as random restarts [128].

However, the approach employed departs from hypermutation as originally presented in that rather than rely upon a trigger, such as a drop in evolutionary fitness, to initiate the increase in mutation factor, instead the hypermutation is triggered randomly. In the scenario investigated, it is indeed not clear what might constitute an effective trigger, since it is not a

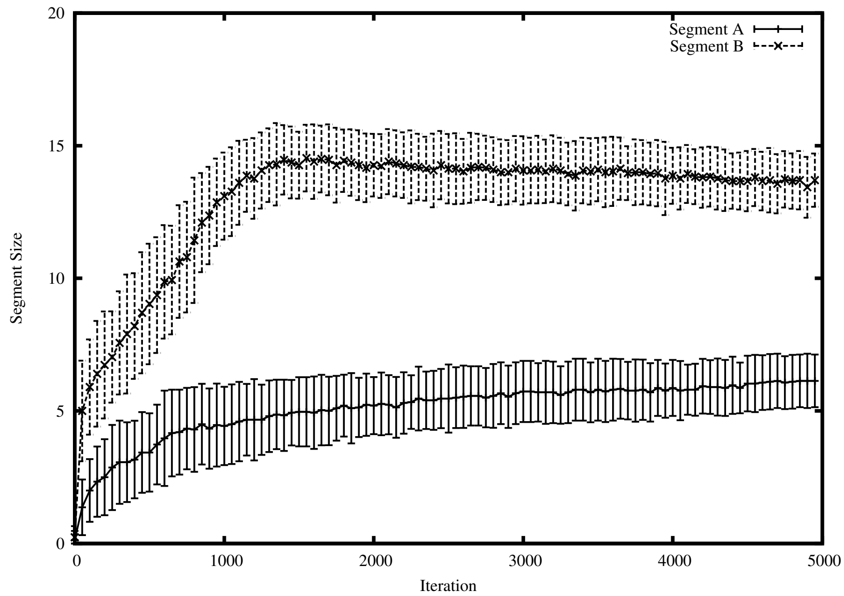

Figure 7.7: Number of sellers in each market segment. The addition of hypermutation allows sellers to continue to move between niches throughout the simulation, without requiring a higher base mutation factor.

drop in fitness or a lack of diversity *per se* which indicates the need for an increased search. Instead, as described in section 7.3, any seller could at any time be in a suboptimal market position without knowledge of this. Therefore, a hypermutation is performed instead of a regular mutation with a very small probability, each time a mutation is made.

In the first hypermutation experiment described, the occurrence of a hypermutation event caused the standard deviation of the Gaussian distribution from which the mutation factor was drawn to increase from 3.0 to 50.0 for that mutation only. Each time a mutation was due to be performed, a hypermutation was instead made with probability 0.01. All else remained the same as the simulation described in section 7.1.

Figure 7.7 shows that despite the initially converged populations being similar to in the version without hypermutation, the addition of random hypermutation allows sellers in segment A to move occasionally to segment B in order to take advantage of the higher payoff available.

As the number of sellers in segment A increases, so the market share served by each seller decreases, and conversely so for segment B. This is reflected in the payoff obtained by the

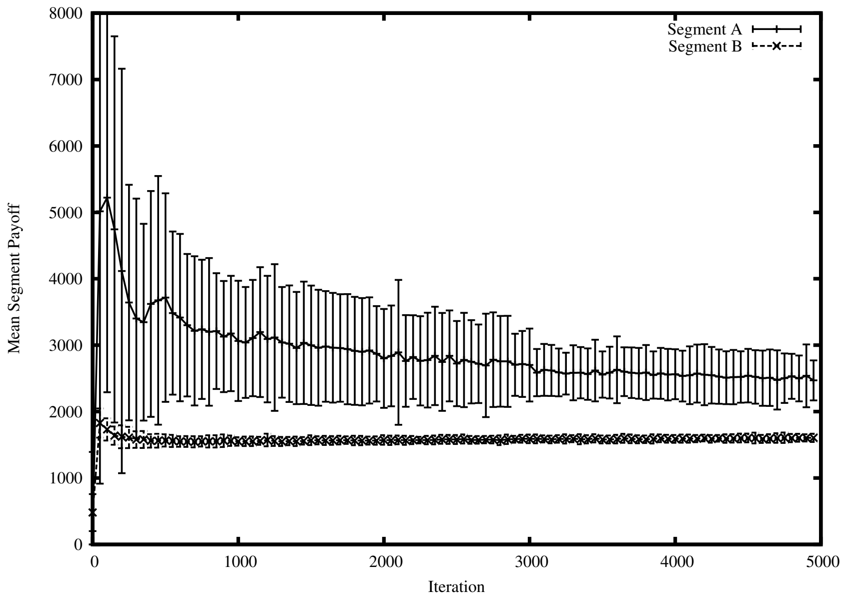

Figure 7.8: Mean and standard deviation seller payoffs for each segment. The addition of hypermutation leads to both more equitable and stable payoffs between the two segments.

sellers, illustrated in figure 7.8. It is clear that ability of even a small number of sellers to move to the more profitable segment A leads to a somewhat more equitable payoff between the segments. However, more importantly, the addition of the hypermutation operator allows the sellers to continue to move between niches throughout the simulation, without requiring a higher base mutation rate and despite their converged populations.

This result in itself is useful, however even given hypermutation, we still do not observe the segments' respective payoffs equalising within any reasonable time, rather drifting together at a slow pace. The rate of equalisation may be increased by increasing the level of hypermutation. In the following set of experiments, the hypermutation operator was modified such that, when triggered, the subsequent n mutations were hypermutations. The probability of a hypermutation occurring remained the same.

Here the mean payoff of the sellers in each segment, calculated between iterations 4000 to 4999 is recorded. This proved to be late enough such that the sellers' populations had completed their initial convergence period, such that any observed population diversity and subsequent movement is due to mutation or hypermutation. Figure 7.9 shows the mean and standard deviation of these values, calculated over for 30 independent simulation runs, plotted

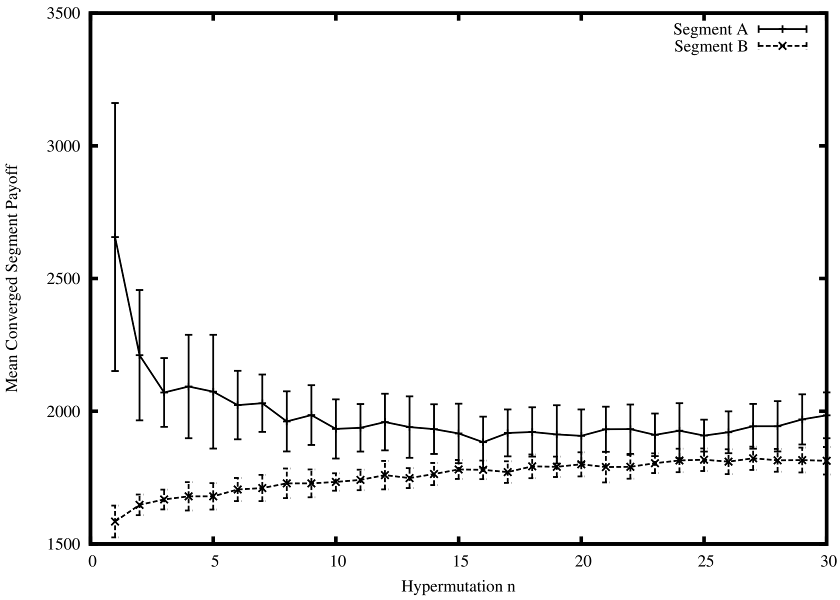

Figure 7.9: Mean sellers' payoff after convergence, by segment, calculated over 30 independent runs, as the number of hypermutations, n, increases. As n grows, the segments' payoffs become more equitable.

as the number of hypermutations per hypermutation occurrence, n varies. Similarly, figure 7.10 shows the mean of the standard deviations between the sellers' payoffs in each run, again calculated over the same 30 independent runs.

As figure 7.9 shows, increasing the number of hypermutations performed, n, leads to increasingly equitable payoffs between the sellers, as an increasing amount of movement is facilitated between the niches, and they become less uneven in size. However, there is a price to pay for this equity; this increase in diversity brought about by more hypermutation leads also to an increasing standard deviation between the sellers' payoffs, as shown in figure 7.10. Since the sellers' populations are more diverse, and the evaluation of each individual contributes towards the total payoff, this result can be expected.

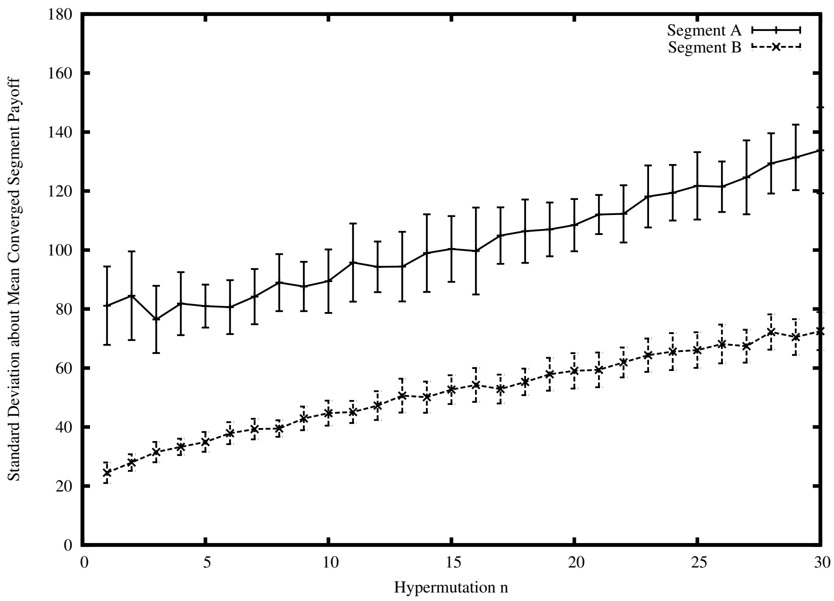

Figure 7.10: Standard deviation between the sellers' payoffs after convergence, by segment, calculated over 30 independent runs, as the number of hypermutations, n, increases. As n grows, the standard deviation between the sellers' payoffs within the segments also increases.

7.5 Expected Seller Behaviour

The discussion in section 7.4 is largely conducted from the perspective of the system as a whole, from a *bird's eye view*, where information about all sellers is available. However, this information is of course not available to the individual sellers, who rely only upon private information. What behaviour then should we expect of self-interested sellers' *evolutionary market agents*, with regard to the diversity of their populations?

Figure 7.11 shows the mean payoff and standard deviation over all sellers in the system, during the converged phase between iterations 4000 and 4999, regardless of their position or payoff. This may be interpreted as an expectation of a seller's payoff during this phase. Again, the number of hypermutations n is varied on the horizontal axis.

Firstly, these results show that the highest expected payoff may be achieved by using no hypermutation at all, when n is zero. However, this comes at a high price in terms of the certainty of the expectation, indicated by the high accompanying standard deviation amongst the sellers. This is because, as was shown, there is a chance of the seller finding a niche which it shares with relatively few competitors. However, this is not guaranteed, and the payoff it receives will be much lower if it finds only the more populated niche. Figures 7.5 and 7.6 illustrated this.

By adopting a small amount of hypermutation (i.e. $n = 1$), a seller can dramatically increase its chance over time of finding the more profitable niche and hence the certainty of the expectation. This comes at a small but significant cost to its expected payoff.

The highest degree of certainty was achieved by sellers in this simulation by performing a small number of hypermutations. Values for n between 2 and 4 give rise to the lowest standard deviations, again with a small penalty in the expected payoff.

Larger number of hypermutations however led to a decrease in the level of certainty with which a seller would obtain the expected payoff, and a decrease in the expected payoff itself. Despite that from the perspective of the system as a whole, this leads to more equitable payoffs and evenly sized segments, this behaviour is hard to motivate from a seller's perspective. Rather, it would be expected that sellers might choose between maximising their expected

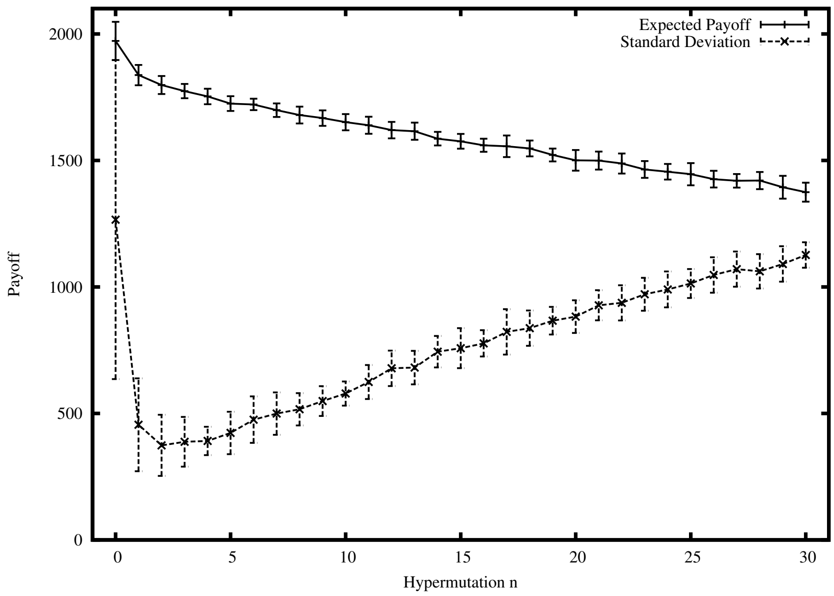

Figure 7.11: Mean of and standard deviation between the sellers' payoffs after convergence, for all sellers, calculated over 30 independent runs, as the number of hypermutations n increases. Rational sellers face a choice between a high expected payoff with a low certainty, and a lower expected payoff with a greater certainty. High levels of hypermutation do not appear beneficial.

payoff, at a cost of certainty, or else sacrifice a small amount of expected payoff in order to buy a greater amount of certainty over their payoff. This is how the diversity dilemma may be characterised from a seller's perspective, and either option leads to the greatest likelihood being an uneven balance of sellers between the segments.

7.6 Conclusions and Discussion

In this chapter the extended multiple attribute market-based approach to resource allocation, as introduced in chapter 6 was further investigated by means of evolutionary simulation. As with the simulations in chapters 4 and 5, resource providers use *evolutionary market agents* to decide the offers for their resources, in a posted offer market mechanism.

It was shown that *evolutionary market agents* are able to position sellers' offers in the market niches identified in chapter 6, providing stable and predictable levels of a quality attribute appropriate to the users. However, a dilemma faced by evolutionary sellers was highlighted, between maintaining a high population diversity to facilitate search and a low diversity to exploit a market position and hence gain a higher payoff. This dilemma was illustrated and characterised using a parameterised hypermutation operator.

We have learnt that the standard *evolutionary market agent* algorithm can lead to sellers self-organising to uneven segment sizes, since populations of offers with low diversity are unlikely to move between niches. As a result, sellers are unable to obtain the optimal payoff. We further learnt that by introducing a randomly triggered hypermutation operator, sellers become more likely to move between niches in order to obtain a higher payoff, making the segment sizes more even. This leads to an increased certainty of achieving the expected payoff. Certainty comes at a cost however, as increased diversity also means that sellers must accept a slightly lower overall expected payoff.

It seems likely that sellers' dilemma is therefore between a higher expected payoff but with low certainty of achieving it, or a slightly lower expected payoff with a higher degree of certainty. The levels of diversity required to achieve truly even segment sizes and fluid movement between segments are unlikely to be realised by self-interested sellers, since in this case both the expected payoff and the certainty are reduced.

The contributions of this chapter are as follows:

- Experimental evidence that sellers with access to only private information, using *evolutionary market agents* are able to position themselves in the market niches identified in chapter 6. The approach is sufficient to generate market segmentation.

- The description of a dilemma facing evolutionary sellers, between maintaining a high population diversity in order to continue to search the attribute space, and a low population diversity with which to exploit their current best known offer.

- The characterisation of the diversity dilemma, through the use of a randomly triggered hypermutation operator, in terms of high expected payoff but with low certainty, and a lower expected payoff with a higher certainty.

- Analysis which suggests that self-interested sellers are unlikely to choose a balance in the diversity tradeoff which would lead to an expectation that the number of sellers in each segment would be representative of its demand.

Since nothing in the implementation of the agents or evolutionary algorithm is specific to the particular way in which the niches in the market were generated, though they were carried out when the multiple peak landscape was generated through *constraint-induced niching*, we expect that the results obtained here will apply equally to other forms of market niching. Nevertheless, it will be important to confirm this experimentally in more complex systems, particularly as Kephart et al. [9] note that multiple peaks may well be the norm in computational markets, as such landscapes may be generated by a wide range of causes.

It is important to note that this chapter does not argue that hypermutation should be a preferred way of framing the described sellers' diversity dilemma or method of tackling it; the operator is used merely as an exploratory tool. Therefore, though the tradeoffs faced by sellers appear to be inherent to the system itself, other diversity mechanisms could lead to potentially more attractive positions within a tradeoff. This will of course depend on the objectives of a particular resource provider, and there is therefore significant scope for numerous other diversity mechanisms to be explored in this context.

Indeed, the hypermutation approach employed here is likely to be suboptimal in a landscape about which sellers have some information. The approach of making large jumps leads more often to a loss of payoff rather than the finding of a more attractive peak. Alternative approaches might therefore include modifying the evolutionary algorithm to make some sort of directed hypermutation, using knowledge obtained about the structure of the payoff landscape in previous iterations. Of course, as with payoff information in the currently used algorithm, it would be important to acknowledge and mitigate against the effects of out of date landscape information, as is done with the stochastic operators discussed in section 4.6.

Of course, these ideas and niching itself are not new in the study and development of evolutionary algorithms [129]. However, since the niching described here is concerned with the proximity of offers, each of which is the result of an independent evolutionary algorithm rather than an individual within a single population, it is unclear what crossover from previous understanding and techniques might be possible. Since each evolutionary algorithm acts independently of the others as it converges to a particular point in the attribute space, no landscape information is shared between algorithms. Any such sharing between algorithms would require private market information to be shared between independent sellers, something which is ruled out by the model and often illegal in real markets.

An important direction for future work will be to investigate the presence of competitive limit-cycle dynamics, as shown to be the case with *bargain hunters* or more stable niching, as seen with *spread buyers* when the buyer population contains a mix of strategies. Section 4.5.3 showed that in the single attribute case, when the proportion of *bargain hunters* increases in a population consisting otherwise of *spread buyers* the stable load balanced outcome obtained with *spread buyers* degrades gracefully. It will be useful to know if similar effects are observed in the multiple attribute case.

CHAPTER 8

CONCLUSIONS AND FUTURE WORK

We know next to nothing about virtually everything. It is not necessary to know the origin of the universe; it is necessary to want to know. Civilization depends not on any particular knowledge, but on the disposition to crave knowledge.

George Will

This thesis has presented a resource allocation problem motivated by an emerging computational paradigm: dynamic, decentralised and massively scalable component-based systems. Making use of the retail-inspired posted offer mechanism, a fully decentralised, evolutionary market-based solution was proposed, which uses competitive co-evolution between self-interested sellers to achieve desired outcome resource allocations. No central or regional coordination is required, nor is any cooperation between agents. The success of this approach was illustrated against a range of resource allocation objectives, including in chapters 3 and 4 an evenly balanced load, in chapter 5 additional stable uneven outcome allocations and in chapters 6 and 7 allocations of resources with multiple quality attributes, appropriate to a heterogeneous population of users.

Crucially, the approach assumes not only a lack of cooperation between agents but also self-interest. Since it accounts for such self-interested utility maximising behaviour, it is more suited to the motivating scenarios than other decentralised approaches. Unlike many market-based systems, the approach requires no central or regional point of control or coordination, such as an auctioneer or specialist. Only private information is available to the agents, as sellers have no knowledge of the size of the market, the last accepted price, the number of competitors or any history. The approach is highly scalable with respect to the number of participating nodes in the system. The time taken for the system to converge to close proximity of the predicted outcome allocation appears to to grow asymptotically to the number of service providing nodes. Additionally, the system has no point which is weaker than any other, and is hence robust to node failure.

Throughout, three buyer behaviour models were considered, hyperrational *bargain hunters*, lightweight *time savers* and risk averse *spread buyers*, each representative of different realistic characteristics. Buyer behaviour models with a step in their decision function were found to lead to unstable, erratic outcome allocations. Those with smooth decision functions however, such as *spread buyers*, led to highly accurate and stable allocations, as well as being robust to small changes in price. It was further shown that this performance degrades gracefully as the proportion of buyers with step functions in the population increases.

Many of the effects described by this thesis were done so through the use of a range of example scenarios and objectives. In each example, the predicted equilibria in the system were analysed and these results compared with those from evolutionary simulations. A major contribution of this, in addition to supporting the claim that particular outcome allocations can be obtained, is to describe and demonstrate a methodology for determining the outcome resource allocation, in the presence of a given configuration of agents. A description was given of how the methodology may easily be extended to cover additional buyer behaviour models not covered here. This generic methodology will be useful in future scenarios and specific embodiments, when further node behaviours need to be considered.

8.1 Summary of Contributions

More specifically, this thesis provides the following contributions:

- A novel approach to the allocation of computational resources in fully decentralised systems using self-interested agents, based on the retail-inspired posted offer market model.

- A methodology for the game theoretic analysis of instances of the approach, enabling the prediction of resulting outcome resource allocations, for a given population of agents.

- A description of how the approach may be applied in order to obtain an evenly balanced load across a decentralised network of resource providing nodes, in the absence of central coordination or cooperation.

- A method for determining suitable parameters for agents in the system in order to bring about additional stable outcome resource allocations. This is achieved transparently to the resource users.

- An extension to the approach, in which resources may be described in terms of multiple quality attributes. Analysis demonstrates that market segmentation in this extended approach can enable resources to be provided with quality of service attributes determined appropriately and proportionally to a heterogeneous population of users.

- A novel agent strategy using evolutionary computation to determine good price and quality attributes for resources, capable of obtaining the predicted outcome allocations using only private information.

- Analysis of the performance of the mechanism when user agents are endowed with a range of different decision making behaviours, and the impact of different buyer decision functions.

While other related approaches to that presented do exist, specifically those which use bilateral bargaining or participation in multiple auctions, the method in this thesis requires less complex reasoning on the part of participating agents. Furthermore, the method proposed is evaluated against its ability to produce desired global outcome resource allocations, such as evenly and unevenly balanced loads. This is not the case of other proposed market-inspired resource allocation techniques. Despite this, the approach taken does not preclude the use of more complex cognitive or otherwise strategic agents. Importantly however, this thesis demonstrates that the simple agents described are sufficient to bring about the desired resource allocations under the assumptions of decentralisation and self-interest.

This is significant, since it will permit designers of systems operating on networks with these characteristics to elicit desired global resource allocations without a high computational cost on individual nodes and without their cooperation. This ensures that the usage of nodes' resources for functional tasks rather than management of the network can be maximised. Furthermore, the scalability and robustness properties of the approach, in particular that of graceful degradation in the presence of disruptive nodes or node failure, ensure that operators of such systems do not need to be as concerned with the impact of such disruptions. This is not the case with systems which require central or regional coordination.

It is important at this stage to also identify weaknesses with the approach and analysis presented. Two assumptions present in the model may be particularly relevant. Firstly there is an assumption throughout this work that network connectivity is uniform, and a viable broadcast capability exists. These assumptions are in fact commonplace in related work, though in embodiments of the model this may not be as straightforward as assumed. Indeed, if message passing is expensive or there exists no way to effectively broadcast offers without relying on other nodes, then additional costs involved in this would need to be added to the model. In particular, if competing nodes are being relied on to forward messages at a cost, then this may affect the dynamics of the market itself, since decisions made in both message passing and resource provision will affect each other.

A further consideration in certain embodiments may be the technical infeasibility or else potential performance hit as a result of breaking a task up into several, indeed many subtasks as is done here. If there are additional costs associated with this, then these should be quantified and also built into the model.

8.2 Future Work

The future directions identified for this work fall into a number of distinct areas. An important next step in making the approach available to system designers will be the embodiment of the studied model in relevant architectures. The model presented in this thesis is purposefully generic and abstract, with the objective of being generally applicable to those classes of system which possess its characteristics: decentralisation, self-interested nodes and generic componentised resources. For this reason, specific technical restrictions and capabilities are not considered here. Therefore, for the approach to be useful in, for example grid systems, compute clusters, cloud and service oriented computing, an important next step is the design or extension of architectures in order to embody the model.

When truly independent self-interest is present, it will be important to consider more complex high level strategies for both buyers and sellers. These will include those able to make use of historical information, as well as those which make use of different approaches to strategy learning, such as the Gjerstad-Dickhaut [116] and Roth-Erev [117] algorithms. The outcomes from this should be compared with existing results and mixed populations also. Crucially, the strategies presented in this thesis act as price takers, despite the effect they have on the market. Higher level strategies may be able to exploit this by moving the market in the short term in order to gain a longer term payoff. As discussed in chapter 3, other buyer behaviours should also be considered. *Sticky buyers* for example, who might prefer to continue using service providers which they have previously used unless there is significant reason to switch, provide a model of potentially highly realistic behaviour. There are of course virtually limitless potential buyer behaviours which could be studied, and this will be made easier by the contribution of a generic methodology for this analysis as presented in this thesis.

It is likely that realistic scenarios will be highly dynamic, where resource providing nodes may be added to or removed from the system during its operation. In addition, the population of resource users may change over time, and there may also be external disturbances. It will be desirable for the system to automatically adjust to such changes, and also to be able to predict how quickly this is achieved. As techniques further develop for the application of evolutionary computation to dynamic optimisation problems, it is likely that this approach will continue to be useful. Nevertheless this is clearly an area where further work would be valuable.

Not least due to future development in the direction of more dynamic scenarios, further design and tuning of the evolutionary algorithm used in the sellers' *evolutionary market agents* can be expected to improve system performance. Despite a small amount of experimental analysis in chapter 3 of the performance of different variants of the algorithm used, this thesis does not claim that this particular evolutionary algorithm is optimal for this purpose; it is merely used to represent the application of an evolutionary process to the problem studied. Further analysis of the algorithm, especially in dynamic environments, will be useful in identifying improvements.

Particularly, market based approaches are widely considered to be highly robust, and the approach presented here, since it it requires no central or regional coordination points, appears to be so. In understanding fully the weaknesses of the evolutionary approach, it will be important to perform more detailed analysis and experimentation in order to quantify this. For example, once equilibrium is reached and a resource allocation stabilised, how robust is this to minor disturbances? This will be particularly relevant in the multiple attribute case, when disturbances in one market segment may impact upon those in another.

Finally, market-based control techniques have been applied to a variety of other real world scenarios other than purely computational ones. Examples of this include air-conditioning control [13] and transport logistics [130] amongst others. The generic nature of the approach presented in this thesis will likely be equally applicable in these scenarios, where the characteristics of self-interest and decentralisation are present.

LIST OF REFERENCES

[1] P. R. Lewis, P. Marrow, and X. Yao, "Evolutionary market agents for resource allocation in decentralised systems," in *Parallel Problem Solving From Nature - PPSN X*, vol. 5199 of *Lecture Notes in Computer Science*, pp. 1071–1080, Springer, 2008.

[2] P. R. Lewis, P. Marrow, and X. Yao, "Evolutionary market agents and heterogeneous service providers: Achieving desired resource allocations," in *Proceedings of the IEEE Congress on Evolutionary Computation*, pp. 904–910, 2009.

[3] P. R. Lewis, P. Marrow, and X. Yao, "Resource allocation in decentralised computational systems: An evolutionary market based approach," *Autonomous Agents and Multi-Agent Systems*, vol. 21, no. 1, pp. 143–171, 2010.

[4] P. R. Lewis, P. Marrow, and X. Yao (named inventors), "Load balancing in a data network." European Patent Application No. 08252953.8, filed September 2008. UK Patent Application No. 0905439.6, filed March 2009. US / Canadian Patent Application No. PCT/GB2009/002135, filed September 2009. Korean Patent Application No. A31728 KRw, filed September 2009. Holder: BT PLC.

[5] I. Foster and C. Kesselman, eds., *The Grid: Blueprint for a New Computing Infrastructure.* Morgan-Kaufman, 1999.

[6] R. Buyya, C. S. Yeo, and S. Venugopal, "Market-Oriented Cloud Computing: Vision, Hype, and Reality for Delivering IT Services as Computing Utilities," in *2008 10th IEEE International Conference on High Performance Computing and Communications*, pp. 5–13, IEEE, 2008.

[7] M. P. Singh and M. N. Huhns, *Service-Oriented Computing: Semantics, Processes, Agents.* Chichester, West Sussex: John Wiley and Sons, 2005.

[8] J. Yang, "Web service componentization," *Communications of the ACM*, vol. 46, no. 10, pp. 35–40, 2003.

[9] J. O. Kephart, J. E. Hanson, and J. Sairamesh, "Price-war dynamics in a free-market economy of software agents," in *Proceedings of the Sixth International Conference on Artificial Life*, (Cambridge, MA, USA), pp. 53–62, MIT Press, 1998.

[10] M. He, N. R. Jennings, and H.-F. Leung, "On agent-mediated electronic commerce," *IEEE Transactions on Knowledge and Data Engineering*, vol. 15, no. 4, pp. 985–1003, 2003.

[11] M. P. Papazoglou, P. Traverso, S. Dustdar, and F. Leymann, "Service-oriented computing: Research roadmap." 2006.

[12] K. Lai, "Markets are dead, long live markets," *ACM SIGecom Exchanges*, vol. 5, no. 4, pp. 1–10, 2005.

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

[14] R. Wolski, J. Plank, J. Brevik, and T. Bryan, "Analyzing market-based resource allocation strategies for the computational grid," *International Journal of High Performance Computing Applications*, vol. 15, no. 3, p. 258, 2001.

[15] D. Cliff and J. Bruten, "Simple bargaining agents for decentralized market-based control," Tech. Rep. HPL-98-17, HP Laboratories, Bristol, UK, 1998.

[16] F. Hsu, *Behind Deep Blue: Building the Computer that Defeated the World Chess Champion*. Princeton University Press, 2002.

[17] A. L. Beberg, D. L. Ensign, G. Jayachandran, S. Khaliq, and V. S. Pande, "Folding@home: lessons from eight years of distributed computing," in *IEEE International Symposium on Parallel and Distributed Processing (IPDPS)*, pp. 1–8, 2009.

[18] E. A. Brewer, "Lessons from giant-scale services," *IEEE Internet Computing*, vol. 5, no. 4, pp. 46–55, 2001.

[19] V. Cardellini, M. Colajanni, and P. S. Yu, "Dynamic load balancing on web server systems," *IEEE Internet Computing*, vol. 3, no. 3, pp. 28–39, 1999.

[20] T. Eymann, M. Reinicke, O. Ardaiz, P. Artigas, L. D. de Cerio F., Freitag, R. Messeguer, L. Navarro, D. Royo, and K. Sanjeevan, "Decentralized vs. centralized economic coordination of resource allocation in grids," *Lecture notes in computer science*, pp. 9–16, 2004.

[21] R. Alfano and G. D. Caprio, "TURBO: An autonomous execution environment with scalability and load balancing features," in *Proceedings of the IEEE Workshop on Distributed Intelligent Systems: Collective Intelligence and its Applications (DIS'06)*, pp. 377–382, 2006.

[22] O. Ardaiz, P. Artigas, T. Eymann, F. Freitag, L. Navarro, and M. Reinicke, "The catallaxy approach for decentralized economic-based allocation in grid resource and service markets," *Applied Intelligence*, vol. 25, no. 2, pp. 131–145, 2006.

[23] V. Ramasubramanian and E. G. Sirer, "Perils of transitive trust in the domain name system," Tech. Rep. TR2005-1994, Cornell University, Ithaca, New York, USA, 2005.

[24] J. Cao, D. P. Spooner, S. A. Jarvis, and G. R. Nudd, "Grid load balancing using intelligent agents," *Future Generation Computer Systems*, vol. 21, no. 1, pp. 135–149, 2005.

[25] J. S. A. Bridgewater, P. O. Boykin, and V. P. Roychowdhury, "Balanced overlay networks (BON): An overlay technology for decentralized load balancing," *IEEE Transactions on Parallel and Distributed Systems*, pp. 1122–1133, 2007.

[26] S. Surana, B. Godfrey, K. Lakshminarayanan, R. Karp, and I. Stoica, "Load balancing in dynamic structured peer-to-peer systems," *Performance Evaluation*, vol. 63, no. 3, pp. 217–240, 2006.

[27] R. Buyya, D. Abramson, J. Giddy, and H. Stockinger, "Economic models for resource management and scheduling in grid computing," *Concurrency and computation: practice and experience*, vol. 14, no. 13-15, pp. 1507–1542, 2002.

[28] P. Rodriguez and E. W. Biersack, "Dynamic parallel access to replicated content in the internet," *IEEE/ACM Transactions on Networking*, vol. 10, no. 4, pp. 455–465, 2002.

[29] R. Chang, M. Guo, and H. Lin, "A multiple parallel download scheme with server throughput and client bandwidth considerations for data grids," *Future Generation Computer Systems*, vol. 24, no. 8, pp. 798–805, 2008.

[30] S. U. Khan and I. Ahmad, "Non-cooperative, semi-cooperative, and cooperative games-based grid resource allocation," in *20th International Parallel and Distributed Processing Symposium, 2006*, p. 10, 2006.

[31] D. K. Begg, S. Fischer, and R. Dornbusch, *Economics*. London, UK: McGraw-Hill, 7th ed., 2002.

[32] M. N. Rothbard, “Free market,” in *The Concise Encyclopedia of Economics*. http://www.econlib.org/.

[33] R. M. Starr, “Commodity money equilibrium in a Walrasian trading post model: An example.” eScholarship Repository, University of California, 2006.

[34] A. Mas-Colell, M. D. Whinston, and J. R. Green, *Micro-Economic Theory*. Oxford: Oxford University Press, 1995.

[35] E. D. Beinhocker, *The Origin of Wealth*. London, UK: Random House Business Books, 2007.

[36] A. Smith, *The Wealth of Nations*. Everyman’s Library, reprint ed., 1991.

[37] Ebay. http://www.ebay.com/.

[38] D. Cliff, “Minimal-intelligence agents for bargaining behaviors in market-based environments,” Tech. Rep. HPL-97-91, HP Labs, Bristol, UK, 1997.

[39] D. P. Friedman and J. Rust, *The Double Auction Market: Institutions, Theories, and Evidence*. Boulder, Colorado, USA: Westview Press, 1993.

[40] S. Phelps, S. D. Parsons, and P. McBurney, “An evolutionary game-theoretic comparision of two double-auction market designs,” in *Proceedings of the 6th Workshop on Agent Mediated Electronic Commerce*, (New York, NY, USA), 2004.

[41] A. Byde, “A comparison between mechanisms for sequential compute resource auctions,” in *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multi-agent Systems*, (Hakodate, Japan), 2006.

[42] E. David, A. Rogers, N. R. Jennings, J. Schiff, S. Kraus, and M. H. Rothkopf, “Optimal design of english auctions with discrete bid levels,” *ACM Transactions on Internet Technology*, vol. 7, no. 2, p. 12, 2007.

[43] N. R. Jennings, P. Faratin, A. R. Lomuscio, S. D. Parsons, C. Sierra, and M. Wooldridge, "Automated negotiation: Prospects, methods and challenges," *International Journal of Group Decision and Negotiation*, vol. 10, no. 2, pp. 199–215, 2001.

[44] A. Byde, "Applying evolutionary game theory to auction mechanism design," in *IEEE International Conference on E-Commerce*, pp. 347–354, 2003.

[45] V. Conitzer and T. Sandholm, "Self-interested automated mechanism design and implications for optimal combinatorial auctions," in *Proceedings of the 5th ACM Conference on Electronic Commerce*, (New York, NY, USA), pp. 132–141, ACM Press, 2004.

[46] D. Cliff, "Explorations in evolutionary design of online auction market mechanisms," *Journal of Electronic Commerce Research and Applications*, vol. 2, no. 2, pp. 162–175, 2003.

[47] H. A. Simon, *Reason in Human Affairs.* Stanford University Press, new ed., 1983.

[48] M. Friedman, *Essays in Positive Economics*, ch. The Methodology of Positive Economics, pp. 3–43. Chicago, USA: University of Chicago Press, 1953.

[49] R. L. Keeney and H. Raiffa, *Decisions with Multiple Objectives: Preferences and Value Tradeoffs.* New York: John Wiley and Sons, 1976.

[50] K. Binmore, *Game Theory: A Very Short Introduction.* Oxford University Press, 2007.

[51] J. V. Neumann and O. Morgenstern, *Theory of Games and Economic Behaviour.* Princeton, NJ, USA: Princeton University Press, 3rd ed., 1953.

[52] J. Nash, "Equilibrium points in n-person games," in *Proceedings of the National Academy of the USA*, vol. 36, pp. 48–49, 1950.

[53] J. M. de Bornier, "The 'cournot-bertrand debate': A historical perspective," *History of Political Economy*, vol. 24, no. 3, pp. 623–654, 1992.

[54] R. Axelrod, *The Evolution of Cooperation.* New York: Basic Books, 1984.

[55] G. Kendall, X. Yao, and S. Y. Chong, *The Iterated Prisoners' Dilemma: 20 Years on.* "World Scientific", 2007.

[56] Y. Cheung, S. Bedingfield, and S. Huxford, "Monitoring and interpreting evolved behaviours in an oligopoly," in *Proceedings of the IEEE International Conference on Evolutionary Computation*, pp. 697–701, 1997.

[57] L. Tesfatsion and K. Judd, *Handbook of computational economics.* Elsevier, 2006.

[58] J. M. Epstein and R. L. Axtell, *Growing Artificial Societies.* Washington, DC, USA: Brookings Institutional Press, 1996.

[59] D. F. Midgley, R. E. Marks, and L. G. Cooper, "Breeding competitive strategies," *Management Science*, vol. 43, no. 3, pp. 257–275, 1997.

[60] M. Wooldridge and N. R. Jennings, "Intelligent agents: Theory and practice," *Knowledge Engineering Review*, vol. 10, no. 2, pp. 115–152, 1995.

[61] A. Sengupta, D. V. Greetham, and M. Spence, "An Evolutionary Model of Brand Competition," in *IEEE Symposium on Artificial Life 2007*, pp. 100–107, IEEE, 2007.

[62] R. E. Marks, "Niche Strategies: The Prisoner's Dilemma Computer Tournaments Revisited." 1989.

[63] R. Axelrod, "The evolution of strategies in the iterated prisoner's dilemma," in *Genetic algorithms and simulated annealing* (L. Davis, ed.), pp. 32–41, London: Pittman, 1987.

[64] J. H. Holland, *Adaptation in natural and artificial systems: An introductory analysis with applications to biology, control, and artificial intelligence.* MIT Press, 1975.

[65] L. J. Fogel, *Artificial Intelligence Through Simulated Evolution.* New York, NY, USA: John Wiley and Sons Inc, 1966.

[66] I. Rechenberg, *Evolutionsstrategien: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution.* Stuttgart, Germany: Frommann-Holzboog, 1973.

[67] X. Yao, "An overview of evolutionary computation," *Chinese Journal of Advanced Software Research*, vol. 3, no. 1, pp. 12–29, 1996.

[68] A. E. Eiben and J. E. Smith, *Introduction to Evolutionary Computing.* Heidelberg: Springer Verlag, 2003.

[69] X. Yao and Y. Xu, "Recent advances in evolutionary computation," *Journal of Computer Science and Technology*, vol. 21, pp. 1–18, January 2006.

[70] X. Yao and P. J. Darwen, "An experimental study of n-person iterated prisoner's dilemma games," in *Selected papers from the AI'93 and AI'94 Workshops on Evolutionary Computation, Process in Evolutionary Computation*, (London, UK), pp. 90–108, Springer-Verlag, 1995.

[71] S. Y. Chong, P. Tino, and X. Yao, "Measuring Generalization Performance in Coevolutionary Learning," *IEEE Transactions on Evolutionary Computation*, vol. 12, no. 4, pp. 479–505, 2008.

[72] S. Y. Chong, P. Tino, and X. Yao, "Relationship Between Generalization and Diversity in Coevolutionary Learning," *IEEE Transactions on Computational Intelligence and AI in Games*, vol. 1, no. 3, pp. 214–232, 2009.

[73] Y. S. Son and R. Baldick, "Hybrid coevolutionary programming for nash equilibrium search in games with local optima," *IEEE Transactions on Evolutionary Computation*, vol. 8, no. 4, pp. 305–315, 2004.

[74] K. De Jong, "Learning with Genetic Algorithms: An Overview," *Machine Learning*, vol. 3, pp. 121–138, October 1988.

[75] J. Maynard Smith, *Evolution and the Theory of Games.* Cambridge University Press, 1982.

[76] S. Phelps, M. Marcinkiewicz, S. D. Parsons, and P. McBurney, "Using population-based search and evolutionary game theory to acquire better response strategies for the double-auction market," in *Proceedings of Trading Agent Design and Analysis Workshop, International Joint Conference on Artificial Intelligence*, (Rochester Hills, MI, USA), pp. 21–27, IJCAI, 2005.

[77] P. Vytelingum, D. Cliff, and N. R. Jennings, "Analysing buyers' and sellers' strategic interactions in marketplaces: An evolutionary game theoretic approach," in *Proceedings of the 9th International Workshop on Agent-Mediated Electronic Commerce*, (Hawaii, USA), pp. 141–154, 2007.

[78] J. Branke, *Evolutionary Optimization in Dynamic Environments.* Kluwer, 2001.

[79] J. Greffenstette, "Genetic algorithms for changing environments," in *Proceedings of the Second International Conference on Parallel Problem solving from nature*, (Amsterdam, The Netherlands), pp. 137–144, Elsevier, 1992.

[80] H. G. Cobb, "An investigation into he use of hypermutation as an adaptive operator in genetic algorithms having continious, time-dependent nonstationary environments," Tech. Rep. Memorandum Report 6760, Navy Research Lab, 1990.

[81] P. Rohlfshagen and J. A. Bullinaria, "Alternative splicing in evolutionary computation: Adaptation in dynamic environments," in *Proceedings of the IEEE Congress on Evolutionary Computation*, pp. 2277–2284, 2006.

[82] J. Branke, "Memory enhanced evolutionary algorithms for changing optimization problems," in *Proceedings of the IEEE Congress on Evolutionary Computation*, vol. 3, pp. 1875–1882, 1999.

[83] A. Gupta, D. O. Stahl, and A. B. Whinston, "The economics of network management," *Communications of the ACM*, vol. 42, pp. 57–63, September 1999.

[84] T. L. Casavant and J. G. Kuhl, "A taxonomy of scheduling in general-purpose distributed computing systems," *IEEE Transactions on Software Engineering*, vol. 14, no. 2, pp. 141–154, 1988.

[85] V. Walia, A. Byde, and D. Cliff, "Evolving market design in zero-intelligence trader markets," in *Proceedings of the IEEE International Conference on E-Commerce*, pp. 157–163, IEEE Computer Society, 2003.

[86] S. Phelps, P. McBurney, and S. Parsons, "Evolutionary mechanism design: a review," *Autonomous Agents and Multi-Agent Systems*, vol. 21, no. 2, pp. 237–264, 2009.

[87] M. Esteva and J. Padget, "Auctions without auctioneers: Distributed auction protocols," in *Lecture Notes in Artificial Intelligence*, vol. 1788, ch. II, pp. 20–28, Berlin, Germany: Springer-Verlag, 2000.

[88] D. Hausheer and B. Stiller, "Peermart: The technology for a distributed auction-based market for peer-to-peer services," in *Proceedings of the IEEE International Conference on Communications*, vol. 3, pp. 1583–1587, 2005.

[89] H. Kikuchi, S. Hotta, K. Abe, and S. Nakanishi, "Distributed auction servers resolving winner and winning bid without revealing privacy of bids," in *Proceedings of the Seventh International Conference on Parallel and Distributed Systems: Workshops*, (Washington, DC, USA), p. 307, IEEE Computer Society, 2000.

[90] J. Niu, K. Cai, S. Parsons, E. Gerding, P. McBurney, T. Moyaux, S. Phelps, and D. Shield, "JCAT: a platform for the TAC market design competition," in *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems: Demo papers*, pp. 1649–1650, International Foundation for Autonomous Agents and Multiagent Systems, 2008.

[91] C. A. Waldspurger, T. Hogg, B. A. Huberman, J. O. Kephart, and S. Stornetta, "Spawn: A distributed computational economy," *IEEE Transactions on Software Engineering*, vol. 18, no. 2, pp. 103–117, 1992.

[92] E. H. Gerding, R. K. Dash, D. C. K. Yuen, and N. R. Jennings, "Optimal bidding strategies for simultaneous vickrey auctions with perfect substitutes," in *Proceedings of the 8th Workshop on Game Theoretic and Decision Theoretic Agents*, (Hakodate, Japan), 2006.

[93] E. H. Gerding, R. K. Dash, D. C. K. Yuen, and N. R. Jennings, "Bidding optimally in concurrent second-price auctions of perfectly substitutable goods," in *Sixth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-07)*, (Honolulu, Hawaii, USA), pp. 267–274, 2007.

[94] T. Eymann, "Co-evolution of bargaining strategies in a decentralized multi-agent system," in *AAAI Fall 2001 Symposium on Negotiation Methods for Autonomous Cooperative Systems*, 2001.

[95] T. Eymann, B. Padovan, and D. Schoder, "The catallaxy as a new paradigm for the design of information systems," in *Proceedings of the 16th World Computer Congress of the International Federation for Information Processing*, 2000.

[96] T. Eymann, M. Reinickke, O. Ardaiz, P. Artigas, F. Freitag, and L. Navarro, "Self-organizing resource allocation for autonomic networks," in *Database and Expert Systems Applications, 2003. Proceedings. 14th International Workshop on*, pp. 656–660, 2003.

[97] E. H. Gerding and J. A. La Poutré, "Bilateral bargaining with multiple opportunities: Knowing your opponent's bargaining position," *IEEE Transactions on Systems, Man and Cybernetics, Part C: Applications and Reviews*, vol. 36, no. 1, pp. 45–55, 2006.

[98] F. Lopes, M. Wooldridge, and A. Q. Novais, "Negotiation among autonomous computational agents: principles , analysis and challenges," *Artificial Intelligence Review*, vol. 2008, no. 29, pp. 1–44, 2008.

[99] A. Chandra, P. S. Oliveto, and X. Yao, "Co-evolution of optimal agents for the alternating offers bargaining game," in *EvoApplications*, vol. 1, pp. 61–70, 2010.

[100] C. R. Plott and V. V. L. Smith, "An experimental examination of two exchange institutions," *The Review of economic studies*, pp. 133–153, 1978.

[101] J. Ketcham, V. L. Smith, and A. W. Williams, "A comparison of posted-offer and double-auction pricing institutions," *The Review of Economic Studies*, vol. 51, no. 4, pp. 595–614, 1984.

[102] S. Jagannathan and K. C. Almeroth, "Price issues in delivering e-content on-demand," *ACM SIGecom Exchanges*, vol. 3, no. 2, pp. 18–27, 2002.

[103] R. Wang, "Auctions versus posted-price selling," *The American Economic Review*, vol. 83, no. 4, pp. 838–851, 1993.

[104] R. B. Myerson and M. A. Satterthwaite, "Efficient mechanisms for bilateral trading," *Journal of Economic Theory*, vol. 29, no. 2, pp. 265–281, 1983.

[105] A. Chavez, A. Moukas, and P. Maes, "Challenger: A multi-agent system for distributed resource allocation," in *Proceedings of the First International Conference on Autonomous Agents*, 1997.

[106] L. Xiao, Y. Zhu, L. M. Ni, and Z. Xu, "GridIS: An incentive-based Grid Scheduling," in *Proceedings of the 19th IEEE International Parallel and Distributed Processing Symposium*, p. 65b, 2005.

[107] K. Kuwabara, T. Ishida, Y. Nishibe, and T. Suda, "An equilibratory market-based approach for distributed resource allocation and its applications to communication network control," in *Market-Based Control: A Paradigm for Distributed Resource Allocation* (S. H. Clearwater, ed.), pp. 53–73, Singapore: World Scientific, 1996.

[108] S. Jagannathan, J. Nayak, K. Almeroth, and M. Hofmann, "On pricing algorithms for batched content delivery systems," *Electronic Commerce Research and Applications*, vol. 1, pp. 264–280, 2002.

[109] S. Phelps, P. McBurney, S. Parsons, and E. Sklar, "Co-evolutionary auction mechanism design: A preliminary report," *Lecture Notes in Ccomputer Science*, pp. 123–142, 2002.

[110] E. J. Anderson and T. D. H. Cau, "Modeling implicit collusion using coevolution," *Operations Research*, vol. 57, no. 2, pp. 439–455, 2009.

[111] T. C. Price, "Using co-evolutionary programming to simulate strategic behaviour in markets," *Journal of Evolutionary Economics*, vol. 7, no. 3, pp. 219–254, 1997.

[112] A. R. Greenwald and J. O. Kephart, "Shopbots and pricebots," *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, vol. 1, pp. 506–511, 1999.

[113] M. Ripeanu and I. Foster, "Peer-to-peer architecture case study: Gnutella network," in *Proceedings of International Conference on Peer-to-peer Computing*, vol. 101, 2001.

[114] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Schenker, "A scalable content-addressable network," in *Proceedings of the 2001 SIGCOMM conference*, vol. 31, pp. 161–172, 2001.

[115] P. Eugster, R. Guerraoui, A. Kermarrec, and L. Massoulie, "Epidemic information dissemination in distributed systems," *Computer*, vol. 37, no. 5, pp. 60–67, 2004.

[116] S. Gjerstad and J. Dickhaut, "Price formation in double auctions," *Games and Economic Behavior*, vol. 22, no. 1, pp. 1–29, 1998.

[117] A. E. Roth and I. Erev, "Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term," *Games and economic behavior*, vol. 8, no. 1, pp. 164–212, 1995.

[118] L. Zeng, B. Benatallah, A. H. H. Ngu, M. Dumas, J. Kalagnanam, and H. Change, "QoS-Aware Middleware for Web Services Composition," *IEEE Transactions on Software Engineering*, vol. 30, pp. 311–327, 2004.

[119] M. Barbuceanu and W.-K. Lo, "Multi-attribute utility theoretic negotiation for electronic commerce," in *Agent-Mediated Electronic Commerce III, Current Issues in Agent-Based Electronic Commerce Systems*, (London, UK), pp. 15–30, Springer-Verlag, 2001.

[120] G. P. Huber, "Multi-attribute utility models: A review of field and field-like studies," *Management Science*, vol. 20, no. 10, pp. 1393–1402, 1974.

[121] G. W. Torrance, M. H. Boyle, and S. P. Horwood, "Application of multi-attribute utility theory to measure social preferences for health states," *Operations Research*, vol. 30, no. 6, pp. 1043–1069, 1982.

[122] P. R. Dickson and J. L. Ginter, "Market Segmentation, Product Differentiation, and Marketing Strategy," *The Journal of Marketing*, vol. 51, no. 2, pp. 1–10, 1987.

[123] S. S. Fatima, M. Wooldridge, and N. R. Jennings, "An agenda-based framework for multi-issue negotiation," *Artificial Intelligence*, vol. 152, pp. 1–45, 2003.

[124] S. S. Fatima, M. Wooldridge, and N. R. Jennings, "Optimal negotiation of multiple issues in incomplete information settings," in *Third International Joint Conference on Autonomous Agents and Multi-agent Systems - Volume 3 (AAMAS'04)*, (New York, NY, USA), pp. 1078–1085, 2004.

[125] W. G. Macready and D. H. Wolpert, "Bandit problems and the exploration/exploitation tradeoff," *IEEE Transactions on Evolutionary Computation*, vol. 2, no. 1, pp. 2–22, 1998.

[126] A. E. Eiben and C. A. Schippers, "On evolutionary exploration and exploitation," *Fundamenta Informaticae*, vol. 35, pp. 1–16, 1998.

[127] J. H. Holland, "Genetic algorithms," *Scientific American*, vol. 267, no. 1, pp. 66–72, 1992.

[128] H. G. Cobb and J. J. Grefenstette, "Genetic algorithms for tracking changing environments," in *Proceedings of the 5th international conference on genetic algorithms*, pp. 523–530, 1993.

[129] P. Darwen and X. Yao, "Every niching method has its niche: Fitness sharing and implicit sharing compared," in *Parallel Problem Solving from Nature—PPSN IV*, p. 398–407, Springer, 1996.

[130] P. Davidsson, L. Henesey, L. Ramstedt, J. Törnquist, and F. Wernstedt, "An analysis of agent-based approaches to transport logistics," *Transportation Research Part C: Emerging Technologies*, vol. 13, no. 4, pp. 255–271, 2005.