

DOCTOR OF PHILOSOPHY

Agent-orientated auction mechanism and
strategy design

Meng Chang

2013

Aston University

Some pages of this thesis may have been removed for copyright restrictions.

If you have discovered material in AURA which is unlawful e.g. breaches copyright, (either yours or that of a third party) or any other law, including but not limited to those relating to patent, trademark, confidentiality, data protection, obscenity, defamation, libel, then please read our [Takedown Policy](#) and [contact the service](#) immediately

Agent-Oriented Auction Mechanism and Strategy Design

A dissertation presented

by

Meng Chang

to

The School of Engineering and Applied Science

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Computational Economics and Multi-Agent System

Aston University

Birmingham, United Kingdom

November 2012

©Meng Chang – November 2012

Meng Chang asserts his moral right to be identified as the author of this thesis.

“This copy of the thesis has been supplied on condition that anyone who consults it is understood to recognise that its copyright rests with its author and that no quotation from the thesis and no information derived from it may be published without proper acknowledgement.”

Agent-Oriented Auction Mechanism and Strategy Design

Abstract

Agent-based technology is playing an increasingly important role in today's economy. Usually a multi-agent system is needed to model an economic system such as a market system, in which heterogeneous trading agents interact with each other autonomously. Two questions often need to be answered regarding such systems: 1) How to design an interacting mechanism that facilitates efficient resource allocation among usually self-interested trading agents? 2) How to design an effective strategy in some specific market mechanisms for an agent to maximise its economic returns? For automated market systems, auction is the most popular mechanism to solve resource allocation problems among their participants. However, auction comes in hundreds of different formats, in which some are better than others in terms of not only the allocative efficiency but also other properties *e.g.*, whether it generates high revenue for the auctioneer, whether it induces stable behaviour of the bidders. In addition, different strategies result in very different performance under the same auction rules. With this background, we are inevitably intrigued to investigate auction mechanism and strategy designs for agent-based economics.

The international Trading Agent Competition (TAC) Ad Auction (AA) competition provides a very useful platform to develop and test agent strategies in Generalised Second Price auction (GSP). AstonTAC, the runner-up of TAC AA 2009, is a successful advertiser agent designed for GSP-based keyword auction. In particular, AstonTAC generates adaptive bid prices according to the Market-based Value Per Click and selects a set of keyword queries with highest expected profit to bid on to maximise its expected profit under the limit of conversion capacity. Through evaluation experiments, we show that AstonTAC performs well and stably not only in the competition but also across a broad range of environments.

The TAC CAT tournament provides an environment for investigating the optimal design of mechanisms for double auction markets. AstonCAT-Plus is the post-tournament version of the specialist developed for CAT 2010. In our experiments, AstonCAT-Plus not only outperforms most specialist agents designed by other institutions but also achieves high allocative efficiencies, transaction success rates and average trader profits. Moreover, we reveal some insights of the CAT: 1) successful markets should maintain a stable and high market share of intra-marginal traders; 2) a specialist's performance is dependent on the distribution of trading strategies.

However, typical double auction models assume trading agents have a fixed trading

direction of either buy or sell. With this limitation they cannot directly reflect the fact that traders in financial markets (the most popular application of double auction) decide their trading directions dynamically. To address this issue, we introduce the Bi-directional Double Auction (BDA) market which is populated by two-way traders. Experiments are conducted under both dynamic and static settings of the continuous BDA market. We find that the allocative efficiency of a continuous BDA market mainly comes from rational selection of trading directions. Furthermore, we introduce a high-performance Kernel trading strategy in the BDA market which uses kernel probability density estimator built on historical transaction data to decide optimal order prices. Kernel trading strategy outperforms some popular intelligent double auction trading strategies including ZIP, GD and RE in the continuous BDA market by making the highest profit in static games and obtaining the best wealth in dynamic games.

Contents

Title Page	i
Abstract	iii
Table of Contents	v
List of Figures	viii
List of Tables	x
Citations to Previously Published Work	xi
Acknowledgments	xii
Dedication	xiv
1 Introduction	1
1.1 GSP Mechanism and Strategy	4
1.2 Double Auction Mechanism Design	5
1.3 Bi-directional Double Auction and Kernel Trading Strategy	6
1.4 Research Contributions	7
1.5 Thesis Structure	8
2 Literature Review	10
2.1 Sponsored Search Mechanism Design	10
2.1.1 Keyword Auction Formats	11
2.1.2 Ranking Mechanism	13
2.1.3 Modelling User Behaviours	15
2.1.4 Analysis of Empirical Data	16
2.2 GSP Strategy	18
2.2.1 Greedy Biding Strategies	18
2.2.2 Budget Constrained Strategies	19
2.2.3 Heuristic Rules	21
2.2.4 Large Scale Optimization	22
2.3 CAT Market Design and Analysis	22
2.3.1 Markets Designed by CAT Entrants	22
2.3.2 Analysis on CAT Specialists	24
2.4 Agent-based Financial Market Simulation	25
2.4.1 General Computational Finance	26

2.4.2	Double Auction Financial Market	28
3	Adaptive Strategy Design for TAC Ad Auction	31
3.1	Sponsored Search and Ad Auction Competition	31
3.2	AstonTAC	34
3.2.1	Bid Price Generator	35
3.2.2	Query Selector	39
3.2.3	Ad Display Selector	41
3.3	Evaluation	42
3.3.1	Game Results and Analysis	42
3.3.2	Controlled Experiments	46
3.4	Conclusion	48
3.5	Future Work	49
4	Mechanism Design of Double Auction	53
4.1	CAT Market Design Tournament	54
4.2	AstonCAT-Plus	56
4.2.1	AstonCAT-Plus Equilibrium Estimation	56
4.2.2	Pricing Strategy	58
4.2.3	Shout Accepting Strategy	58
4.2.4	Clearing Strategy	60
4.2.5	Charging Strategy	64
4.3	Evaluation	66
4.3.1	Heterogeneous Games	67
4.3.2	Head-to-Head Games	75
4.3.3	Trader Distribution Analysis	80
4.3.4	Trading Strategy Preferences	84
4.4	Conclusion	86
5	Bi-directional Double Auction and Kernel Trading Strategy	88
5.1	Market Mechanism	88
5.1.1	BDA Market Architecture	88
5.1.2	Agent Mixing Manager	91
5.1.3	Trading Direction Algorithms	91
5.1.4	News System	95
5.1.5	Trading Strategies	95
5.2	Kernel Trading Strategy	97
5.3	Controlled Experiments	98
5.3.1	BDA Market Time Series	99
5.3.2	Static Market Efficiency	101
5.3.3	Performance of Kernel Strategy	105
5.4	Conclusion	108

6 Conclusion and Future Work	111
Bibliography	115
A Parameter Settings of AstonTAC and TAC AA competition	130
B List of Symbols Used in Chapter 3	132
C Parameter Settings and Symbol Denotation of AstonCAT-Plus	134
D Table of Experiment Outcomes	137

List of Figures

3.1	Daily activities cycle in TAC AA.	32
3.2	Architecture of AstonTAC.	34
3.3	Average revenue and profit of top three agents.	43
3.4	(a) Average CTR (b) Average CPC (c) CPC/CTR of AstonTAC, Schlemazl and Tactex.	43
3.5	Click through rate against capacity.	44
3.6	Distinct bidding patterns shown by AstonTAC (wave pattern on the left) and TacTex (jigsaw pattern on the right).	45
3.7	Different variances of top three agents in terms of daily conversion.	45
3.8	Winning margin and frequency comparison for experimental game 690-769.	47
3.9	Different CTR trend against change of distribution capacity (Experiment C).	48
4.1	AstonCAT-Plus architecture.	56
4.2	Demonstration of AstonCAT-Plus's shout accepting strategy.	58
4.3	Hierarchical market-adaptive stabilized charging strategy.	64
4.4	Profit fees of AstonCAT-Plus in comparison with its market shares (Data from heterogeneous game iteration 8).	66
4.5	Score comparison for heterogeneous games (AstonCAT is not in heterogeneous games).	68
4.6	Average Trader Profit Per Transaction. Data are from the 8th iteration of the heterogeneous games.	70
4.7	Average market share.	71
4.8	Average transaction number (goods traded) per trader.	71
4.9	Comparison between TPT-CDA and CDA (only) clearing strategy.	72
4.10	Allocative Efficiency of each agent in a typical heterogeneous game.	74
4.11	Convergence coefficient of each agent in a typical heterogeneous game.	76
4.12	Daily market shares in head-to-head games.	79
4.13	Distributions of universal traders, intra-marginal traders and equilibrium trader profits in a typical heterogeneous game.	80
4.14	Mean and standard deviation of percentage of intra-marginal trader to total number of traders.	81
4.15	Trader side-balance rate λ of each specialist.	83

4.16	Average daily score of each specialist in Type I games.	85
4.17	Average daily score of each specialist in Type II games.	85
4.18	Average market share of each specialist in Type II games.	86
5.1	Architecture of BDA market.	89
5.2	Daily activity flow of BDA market.	90
5.3	Building kernel density estimators.	94
5.4	Mesh of private valuation update ratio function eliminating the stochastic part.	96
5.5	Daily closing prices in iteration 6 and corresponding news series on the same time scale.	99
5.6	Decaying auto-correlation of absolute daily log returns in BDA market. . .	101
5.7	Convergence of transaction prices in static games.	102
5.8	Sorted market efficiencies of 100 static games.	102
5.9	Sorted market efficiencies categorised by uncertainty degree distributions of the valuation.	103
5.10	Average market efficiencies of static game in regard to uncertainty distributions of the valuation.	103
5.11	Daily cumulative market efficiency in a game.	104
5.12	Average market efficiencies and deviations regarding trading strategies. . .	105
5.13	Average profit of each trader group in static heterogeneous game (100 iterations).	106
5.14	Average traded quantity of each trader group in static heterogeneous game (100 iterations).	106
5.15	Mean of the average profit per unit of each trader group in static heterogeneous game (100 iterations).	107
5.16	Proportional comparison of profit, traded quantity and unit profit between <i>Dual</i> ZIC trader group and <i>Bi</i> ZIC trader group (100 iterations).	107
5.17	Final wealth each trader group possess in dynamic heterogeneous games. .	109
5.18	Aggregate group wealth of dynamic heterogeneous games (20 iterations). .	109

List of Tables

3.1	The algorithm of query selector.	39
3.2	Settings of post-tournament ad auction games.	46
4.1	Pseudo code of AstonCAT-Plus' market clearing strategy.	62
4.2	Specialists used in controlled experiments.	67
4.3	Transaction success rates summary for heterogeneous games.	69
4.4	Mean and standard deviation of average profit per trader of across 10 heterogeneous game iterations.	72
4.5	Summary of allocative efficiency and convergence coefficient in heterogeneous games.	73
4.6	The results of head-to-head games.	77
4.7	Trader population settings for strategy distribution games.	84
5.1	BDA market experiment configuration details.	100
5.2	Market average efficiency and standard deviation calculated based on the data of 100 static game iterations.	101
A.1	Parameter Settings of AstonTAC and TAC AA competition.	131
B.1	List of symbols used in Chapter 3.	133
C.1	Parameter settings and symbol denotation of AstonCAT-Plus - Part 1. . . .	135
C.2	Parameter settings and symbol denotation of AstonCAT-Plus - Part 2. . . .	136
D.1	CTR and profit of TacTex and AstonTAC in controlled experiment C of Chapter 3.	138

Citations to Previously Published Work

Most of Chapter 3 has been published as:

“Designing a Successful Adaptive Agent for TAC Ad Auction”, Meng Chang, Minghua He and Xudong Luo, Proceeding of the 19th European Conference on Artificial Intelligence (ECAI), Lisbon, 136-139, 2010

Large portions of Chapter 4 have appeared in the following paper:

“AstonCAT-Plus: An Efficient Specialist for the TAC Market Design Tournament”, Meng Chang, Minghua He and Xudong Luo, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, 146-152, 2011

Finally, the essential content of Chapter 5 is discussed by paper:

“Bi-directional Double Auction for Financial Market Simulation”, Meng Chang, Xudong Luo, Aniko Ekart, Minghua He and Shichao Zhang, Proceedings of the 12th International Conference on Autonomous Agents and Multi-agent Systems (AAMAS 2013).

Acknowledgments

Completing this doctoral work has been a wonderful and often overwhelming experience. It is hard to know whether it has been grappling with the agent-based computational economics itself which has been the real learning experience, or grappling with how to write a paper, give a talk, work in a group, lab demonstration, recover a crashed computer, stay up until the birds start singing, and work continuously while fighting with cold and flu.

I have been very privileged to have undoubtedly the most intuitive, smart and supportive supervisor anyone could ask for, namely Minghua He. I have been stimulated and excited by her constant flow of good ideas. Minghua has an ability to guide me in the right direction of research and point out the important issues I miss to consider in a project. When I am stuck with any problem, she always shows me the door to the solution rather than the solution itself. Through this process, my research ability is significantly improved after four years of practice. Minghua has a warm heart and open mind. She is always there to encourage me when I am frustrated by failures and support me when I come up with new ideas.

Throughout my four years of PhD study, I have been helped and supported by so many people who I am so grateful to. First, I have gain essential knowledge in machine learning and multi-agent system through the lectures of Prof. Ian Nabney and Dr. Maria Chli. In my BDA project, Dr. Aniko Ekart has been brain storming with me and corrected many mistakes or potential misunderstanding points in my proposal of the system. Her valuable suggestions have helped me to perfect the model design. Dr. George Vogiatzis has been working together with me in lab teaching of 3D animation. His paper about a full probability model has given me inspiration of solving economical problems using probability based models. As both friend and mentor, his trust and recommendation have enhanced my faith in myself in terms of creating something good for the world. In addition, I would like to say, without recommendation of my MSc supervisor Dr Evor Hines in Warwick University, I would not have the opportunity to demonstrate my potential and ability in academic research in computer science.

Special thanks must be given to Prof. Xudong Luo who has been my guiding light in the process of paper writing and publishing over the last four years. He generously passes his experience in composing professional publications to me, which helps me to avoid many obstacles that new academic writers often come across and save immense amount of time. Without his diligent help and inspiring guidance, I simply cannot imagine completing so many quality papers within my PhD when the majority of my time has to be spent on running experiments and collecting data.

Sandra Smith, our research student administrator, is surely the kindest, coolest and most

witty person one could possibly hope to spend a lunch-break gossiping with. Her advice on my recent dog bite accident has given me great peace of mind. Her prompt and accurate working style is highly appreciated.

There are countless others who have been there for me throughout my time as a PhD student. From the world of agent-based e-commerce, there are Michael Winsper, Micheael Garlic, Rodrigue Talla kuate, and so on. From other worlds in computer science, there are Dr. Hai Wang, Dr. Dan Cornford, Mitchell Eccles, James Keane, Chunhui Li and Darian Brookes.

My fascination of becoming a scientist is undoubtedly due to the influence of my father Mr. Shaoyu Chang and my mother Mrs. Yuzhi Zhao. From very young, they started to stimulate my interests in science. My mother is a maths teacher in a primary school. Although she now cannot even understand the equations and theorems I use in my research, I should never forget my mathematical mind was enlightened by her initially. My father did not have the chance to receive high level education when he was in my age. But he deeply knows the importance of knowledge and spares no effort to support my study in my interested subjects. I still remember when I was in Xiamen University, I was one of the first a few students who are lucky to own an expensive personal computer, which becomes my strongest motivation to explore the world of computing.

The last but not least person I have to thank is my wife Chenrong Lin. She gave me immense amount support during not only my four years of PhD at Aston University but the last eight years while we are together in the UK from every respect imaginable. No matter through lows or highs, I can always be calm when she is beside me. Her selfless tendance on me is like a fountain that moisturises my heart and soul. In terms of my research, her constant encouragement and constructive advices give me the endless power to face even the toughest challenges.

*Dedicated to
my father Shaoyu Chang,
my mother Yuzhi Zhao,
my wife Chenrong Lin.*

Chapter 1

Introduction

In recent years, electronic commerce (e-commerce) has become an important tool for businesses worldwide in terms of not only boosting sales but also engaging customers [52]. It offers great opportunities to significantly improve or permanently change the way that businesses interact with both their customers and suppliers. In fact, new commodities such as search engine keywords, e-money [65] are created in e-commerce. The sponsored search helps businesses gain exposure to their potential customers through displaying their adverts and website links when users search keywords that are related to their businesses. E-money provides Internet companies a convenient way to charge their customers for the services they provide. Furthermore, fast development of computers and Internet enables many new business models such as group buying [14] in which the Internet helps to locate individual buyers of the same product or service and the group buying company organise them into a buying group in order to obtain the special discount for the simultaneous purchase of multiple items. In business-to-business world, the integration of supply chain significantly improves operational efficiency and reduces operational costs. Such systems cannot be implemented without intervention of computing technologies.

As e-commerce is increasingly assuming pivotal role in almost all kinds of businesses, how to act quickly, wisely and strategically in the new environment is realised as a new challenge. To address this challenge and harness the full potential of e-commerce, we believe that a new model of software is needed. This model is based upon the concept of *software agents* – software entities that act on behalf of their owner in an autonomous fashion in order to achieve their objectives [81]. Compared with human beings, the software agent is able to collect and analyse much larger amounts of information in much shorter time and process thousands of requests simultaneously. Agents are considered intelligent if they can make strategic decisions or take optimal actions in some environment automat-

ically. When two or more cooperative organisations use agents, their agents should be able to autonomously communicate, coordinate and collaborate with one another and ultimately form coalitions in order to achieve their common goals.

In agent-based computational economics, economic processes are modelled as dynamic systems of interacting agents [184]. If multiple interacting intelligent agents are involved, such a system is called multi-agent system [183]. In the design of agent-based systems, there are two key aspects to be considered. One is concerned with the agent itself and the other is concerned with the relationship between different agents. First, an agent needs to have some strategies to guide its behaviours in order to achieve its design objectives within a specific environment. To form effective strategies, the agent usually needs certain level of intelligence in terms of perceiving and reacting to the change of its environment, learning and inferring from historical data, and predicting the trend of future [82]. Second, agents need to interact with one another. Therefore, the interacting mechanism needs to be carefully designed. In economic systems, agents interactions often result in transactions and resource allocations. With self-interested agents, the most widely studied and employed interaction mechanisms are auctions - institutions where goods are sold by the process of making bids and allocating goods according to competition [71]. Auctions are prevalent because they are extremely efficient and effective methods of allocating goods or services [185]. In economic theory, an auction may refer to any mechanism or set of trading rules for exchange. Hence, auctions come in many different forms, each with their own rules and ensuing properties.

There are four primary types of auctions:

- **English auction**, also known as an open ascending price auction [71]. In English auction, the auctioneer starts with a reserve price and solicits successively higher public bids from the bidders until no one will increase the bid and the last bidder is the winner who pays for his/her bidding price.
- **Dutch auction**, also known as an open descending price auction. In Dutch auction, the auctioneer begins with a high asking price which is lowered until some participant is willing to accept the auctioneer's price [118]. The winning bidder pays the last announced price.
- **First-price sealed auction**, in which all bidders simultaneously submit sealed bids so that no bidder knows the bid of any other participant [118]. The highest bidder wins the item and pays the price he/she submitted.

- **Second-price sealed-bid auction**, first described by William Vickrey [173] where bidders submit written bids without knowing the bid of the other people in the auction, and in which the highest bidder wins, but the price paid is the second-highest bid.

Other popular auctions include bidding fee auction [136], unique bid auction [148], reverse auction [79], and combinatorial auction [142].

Given the variety of auction protocols, bidding strategies must be tailored to the type of the auction in order to be effective or just working. Therefore, agent-oriented auction mechanism and strategy design become an interesting problem to us. Our research is mainly related to two specific types of auctions:

- **Generalised Second Price auction (GSP)**, in which bids are sorted in descending order, the highest bidder wins the top slot, the second highest bidder wins the second highest slot and so on. Each bidder pays the bid of the next highest bidder if he/she wins a slot [50]. It is used mainly in the context of keyword auctions, where sponsored search slots are sold on an auction basis. This technology is employed by most search engines including Google, Bing, Yahoo and Baidu in order to fairly and efficiently set prices for the right to use their advert slots available on a search result page of any keyword or keywords combination. We survey the mechanism and bidding strategy design issues of GSP in the context of sponsored search (see Chapter 2) and present an effective advertiser strategy that shows superior performance in the TAC AA competition (see Chapter 3).
- **Double Auction (DA)**, in which multiple sellers submit asks and multiple buyers submit bids continuously and simultaneously, and the auctioneer choose a price p to clear the market [71]. In terms of the time delay of clearing the market after finding match(es) between asks and bids, there are major variations of DA: Continuous Double Auction (CDA) which clears the market as soon as a match is found and Clearing House (CH) which clears that market at the end of fixed periods. This auction, often appearing in the form of the combination of CDA and CH, is the dominant exchange institution adopted by financial markets for trading equities, currencies, derivatives etcetera. Especially, the CDA has been the principle trading system in equity markets for more than 140 years under open-outcry floor trading [188]. Due to its high efficiency in resource allocations, it is also used to solve many other allocation problems including bandwidth allocation and electricity distribution. With respect to shout accepting, matching and charging rules etcetera, it has other interesting alterations than

just CDA and CH. Our research about DA is first concentrating on the automatic market design based on the framework of TAC CAT (see Chapter 4). Then, we develop a Bi-directional Double Auction (BDA) model for financial market simulation, in which a high-performance trading strategy based on probability density estimations is proposed and evaluated (see Chapter 5).

In the following, we introduce the three topics we are going to discuss in this thesis: 1) GSP mechanism and strategy, 2) Double auction mechanism design and 3) Bi-directional double auction and Kernel trading strategy.

1.1 GSP Mechanism and Strategy

In recent years, sponsored search [54, 97] has become the indispensable source of revenue for Internet search engine companies like Google, Yahoo and Bing. Instead of showing the same advertisement to every user, it enables companies to promote their products to targeted groups of consumers based on their search queries [151]. Moreover, setting a price for the advertisement is through Generalised Second Price auction, which has many advantages over the conventional negotiation between the seller and the buyer such as price efficiency. Since GSP is mainly used to sell keywords, it is often referred as the keyword auction.

From agent point of view, keyword auction can be modelled as a multi-agent system containing three types of agent - the user, the advertiser and the publisher (the search engine). The user wants the most relevant ads to appear in the most prominent place so they can spend the smallest amount of time on finding what they want. The advertiser wants to maximise his return of investment. The publisher wants to satisfy both the user and advertiser with relevant ad, fair ranking and charging, stable amount of traffic, etcetera., to make sure they will come back and reuse their services. More importantly, search engines obviously want to maximise their long-term return in revenue through valuable services to both parties.

Since keyword auction has become the dominant method that search engines adopt to sell their advertisement services, the advertisers face one inevitable question: how can they formulate an effective bidding strategy under the rules of GSP. This issue has also attracted attentions of computer scientists. In order to explore the solutions, University of Michigan developed Trading Agent Competition Ad Auction (TAC AA) [84], in which participants play the role of advertiser in a simulated keyword auction environment and compete with

each other in order to maximise their own profit. In the first TAC AA tournament held in 2009, we designed the software agent AstonTAC, which was ranked the 2nd and made the highest revenue in the competition. Our strategy allows us to extract the Market-based Value Per Click (MVPC) of every keyword and perform adaptive bidding along with the changes of MVPC. Our paper “Designing a Successful Adaptive Agent for TAC Ad Auction” [28] describes our findings in regard to GSP strategy formulation in a multi-keywords environment.

1.2 Double Auction Mechanism Design

Double Auction market is a market in which multiple buyers compete to purchase many items that are simultaneously offered for sale by multiple sellers competitively [72]. This mechanism has dominated today’s financial instruments exchanges (*e.g.*, the New York Stock Exchange and the London Stock Exchange) for its high allocative efficiency and simplicity in implementation. As economy and technologies evolve, the burgeoning on-line trading system and electronic marketplaces have offered traders more freedom than ever to trade their securities across the world. Given this, one double auction market has to face competitions from other similar markets running concurrently around the world [155]. So, to design an efficient market in such a competitive environment, we need to address the following issues well: 1) How can such a market attract and keep traders in such a competitive environment? 2) How can such a market maximise its own profit by charging as much fees as possible without driving traders away for the sake of overcharging? And 3) how to facilitate the most efficient allocation of resource inside such a market?

It is too complex to analyse double auctions, particularly continuous double auctions, theoretically [59]. Thus, researchers turned to do their experimental analysis by developing software agents that can bid for goods or services on behalf of their human owners [71, 114]. Since 2007, market design tournament (also called CAT) was introduced into the International Trading Agent Competition¹ to simulate the competitive environment of multiple double auction markets and enable automated market mechanism design. It aims to seek answers to the above questions.

AstonCAT is a double auction mechanism designed in the CAT tournament. Inspired by its soaring improvement on performance in Game 3 of CAT-2010 (rank advanced by 4 places), we developed a post-tournament version called AstonCAT-Plus, which signif-

¹<http://www.sics.se/tac/>

icantly outperforms its predecessor and achieves the highest transaction success rate, allocative efficiency and average trader profit among all the specialists of double auction in our controlled experiments.

More specifically, AstonCAT-Plus uses a market equilibrium framework - both the shout accepting thresholds and the estimated equilibrium prices are based on the estimated market equilibrium price. We choose the equilibrium framework because a market reaches the best allocative efficiency if it is always cleared on its equilibrium price and an equilibrium clearing market encourages truthful bidding behaviour of the traders. It adapts the shout accepting thresholds against the change of market conditions and clears the market based on the profit per trader. This allows the shout engine to match more profitable bid-ask pairs. Moreover, we treat matched shouts differently if one of the shouts in the matched pair is an extra-marginal shout identified by comparing the shout price to the estimated equilibrium price. Further, our charging strategy constrains the fees to small ranges while adapting it to market share related criteria.

1.3 Bi-directional Double Auction and Kernel Trading Strategy

Nowadays, computer scientists are increasingly involved in building market systems [172] that often employ double auction mechanism because of its high efficiency of resource allocation [180, 146]. Consequently, the CAT competition [22] has been developed to investigate the optimal design of a double auction market. Although many effective alterations of double auction market have been proposed since the competition [132], researchers typically only deal with one-way traders *i.e.*, traders are either buyers or sellers but not both [140]. It is well-known that the dominant application of double auction institution is the financial market, in which traders are usually sellers and buyers simultaneously. Hence, we consider necessary to introduce a Bi-directional Double Auction (BDA) model, in which the trading activity of every individual trader can be bi-directional. To complete a dynamic financial market simulation, we also introduce a news system to enable traders to update their private valuations along with the change of the environment.

In the BDA market, the selection of trading direction is through trading direction algorithms, for which we introduce *Dual* and *Bi*. Once the decision is made, the order price is determined by a trader's trading strategy. Besides implementing some of the most popular double auction trading strategies (see Section 5.1.5), we develop a new trading strategy

called *Kernel* based on probability density estimations, which significantly outperforms all other strategies in our experiments.

1.4 Research Contributions

The work described in this thesis makes a number of important contributions to the state of the art in the area of auction mechanism and trading strategy design and in the realm of agent-based computational economics.

- We design a successful bidding strategy for GSP-based keyword auction in the environment of TAC Ad Auction competition. Specifically,
 1. We introduce a novel concept called Market-based Value Per Click (MVPC). Based on MVPC, we submit bids that win appropriate ad slots at reasonable costs.
 2. Given limited conversion capacity, we propose an effective query selection algorithm which estimates the realistic maximum number of conversion allowed by the advertiser's distribution capacity and sort queries by their profitability. The algorithm only selects the most profitable queries to bid on so that the demands of conversion quantity and high conversion value are balanced.
- We design a double auction market mechanism which demonstrates strong and stable performance on TAC CAT platform. Specifically,
 1. We develop a novel clearing strategy which clears the market considering both the quality and quantity of transactions.
 2. We introduce a new shout accepting policy that adapts the accepting thresholds dynamically according to the change of market condition.
 3. We propose an algorithm that combines the long-term and short-term transactional data to effectively estimate local market equilibrium price.
 4. We introduce a rule-based hierarchical charging strategy, which effectively balances maintaining high market share and generating high revenue in a competitive multi-market environment.
- We analyse the CAT tournament to reveal essential features of successful CAT specialists. For example,

1. CAT specialists' performance is affected by the distribution of trading strategies. No specialist is absolutely superior to others if the current dominant trading strategy is not preferred by this specialist.
 2. A successful market should maintain a stable and high market share, especially the share of intra-marginal traders.
 3. A successful specialist usually features a balanced profile between buyers and sellers.
- We introduce the Bi-directional Double Auction (BDA) for financial market simulation in which traders determine their trading directions dynamically as what they do in real financial markets. For the BDA market, we create *Dual* and *Bi* trading direction algorithms to model the traders' behaviours in terms of changing trading directions when trading financial instruments. Furthermore, we reveal properties of the static continuous BDA Market:
 1. The market allocative efficiency largely comes from traders' rational choices of trading directions.
 2. With incentive-compatible trading direction algorithms, the more intelligent the trading strategies, the less efficient the market.
 3. The market is more efficient and stable if traders are more confident with their private valuations of the traded asset.
 - Last but not the least, we develop *Kernel* trading strategy in the continuous BDA market, which shows superior performance when competing with some popular existing intelligent CDA strategies adapted for the BDA market. This strategy, for the first time, uses kernel probability density estimation technology to calculate the transaction probability of future shouts and submit shouts to maximise the expected profit.

1.5 Thesis Structure

The thesis is structured as follows:

- Chapter 2 summarises the related work to our research in generalised second price auction and double auction. For GSP, we first survey existing work in regard to

sponsored search mechanism design. Then, we focus on the bidding strategy problem faced by the advertiser. For DA, we present the ideas of other successful CAT tournament entrants and post-tournament analytical findings regarding double auction mechanisms. Finally, we design a bi-directional double auction financial market simulation and discuss various ways of implementing artificial financial markets in addition to a general introduction of agent-based computational finance which is the background of this work. Meanwhile, some classical design of double auction trading strategies are described because our new *Kernel* strategy is inspired by the study of these existing ones.

- Chapter 3 specifically describes the design, implementation and evaluation of the TAC AA agent AstonTAC. We emphasise the importance of understanding the feature of the commodity and evaluation of its market value against a dynamic background. We also point out, with limited budget and conversion capacity, selecting the right keywords to invest is as important as employing an effective strategy on the bidding of any individual keyword.
- Chapter 4 concentrates on the TAC CAT competition for which AstonCAT-Plus designed. From the aspects of equilibrium estimation, accepting policy, clearing policy and charging policy, we introduce our innovative ideas that make AstonCAT-Plus one of the most efficient and attractive double auction mechanisms in this competitive multi-market environment. In addition, we reveal some insights about the CAT competition through trader analysis and trading strategy distribution games.
- Chapter 5 introduces the bi-directional double auction market. In details, we describe *Dual* and *Bi* trading direction algorithms and introduce a high-performance *Kernel* trading strategy. Through experiments, we analyse endogenously generated time-series and show that BDA is a valid model for financial market simulation. Furthermore, we investigate the allocative efficiency of the static continuous BDA market and evaluate the performance of *Kernel* trading strategy with heterogeneous trading games in both static and dynamic settings.
- Chapter 6 recaps the main contributions of this thesis and highlights the key open problems that need to be addressed in both GSP and DA markets.

Chapter 2

Literature Review

In this chapter, we review the literature related to our research. Based on the three separate projects involved in our research and a short survey conducted on sponsored search, we divide our review into four parts: (i) Sponsored search mechanism design; (ii) GSP strategy; (iii) CAT market design and analysis; (iv) Agent-based financial market simulation.

2.1 Sponsored Search Mechanism Design

In sponsored search, advertisers enter an amount that they are willing to pay for showing their ads alongside the algorithmic results of a keyword. The search engine sorts advertisers' bids by some sort of ranking mechanism and awards positions to advertisers' ads. When the user searches the keyword targeted by participated advertisers, he will see their ads as well as natural results. If he clicks on one or more ads, the corresponding advertiser will have to pay the search engine for the clicks by a price determined by the auction. Advertisers are normally ranked by a score which may not only relate to the bid price but also the corresponding ad's relevance, quality, and so on. In this scenario, the search engine is responsible of designing a practical and efficient auction mechanism that attracts advertisers and satisfies its users by automatically providing helpful sponsored link in a none-interruptive way. According to a specific auction mechanism, the advertiser needs to formulate a bidding strategy in order to maximise its expected utility under some budget constraint typically.

Jordan *et al.* identify two main streams in the research of sponsored search [84]: (i) mechanism design problem faced by search publishers, (ii) strategic problem faced by advertisers. Without comprehensive understanding of the market mechanism, it is impossible to create a successful strategy in it. Therefore, the survey of the sponsored search mech-

anism is very necessary. The mechanism design has to be based on certain assumption of user behaviour model such as viewing, clicking and conversion behaviours. So what are the appropriate models for them? Usually independent ad effect is assumed, but is it true? Have we missed out important externalities that actually exist in the real business? Generalised GSP is the most popular mechanism used by search engines. Why is it selected? What are its merits and drawbacks? What other mechanisms have been proposed and why are they not popular? With these questions and more, we discuss literature on the design of sponsored search mechanism.

2.1.1 Keyword Auction Formats

In sponsored search, keyword auctions are employed to sell ad slots automatically. Therefore, conventional auction formats are generalised to create keyword auction mechanisms. Many formats have been tried and GSP [50, 73, 122] finally becomes the industrial standard because it makes the market more user friendly and less susceptible to gaming. Big search engines like Google and Yahoo adopted it quickly after recognising its advantages.

GSP, as its name implies, is conceived as a natural extension of second price seal-bid auction, also called Vickrey auction [107]. In details, GSP auction has two major formats: rank-by-bid GSP and rank-by-revenue GSP. In rank-by-bid GSP auction, the advertiser in position i pays a price per click equal to the bid of an advertiser in position $i + 1$. In rank-by-revenue GSP auction, a weight w_s is associated with bidder s . If bidder s bids b_s , his corresponding score is $w_s b_s$ which determines his rank. Without loss of generality, let us assume bidder s receives position s , what he pays is the minimum payment p to retain his position $p = \frac{w_{s+1} b_{s+1}}{w_s}$. Although conventional second price sealed auction has the property of truthful bidding, GSP is proved no longer truthful after generalisation [50].

However, besides GSP, there is not lack of other popular formats that have been used, proposed or analysed.

Generalised First-Price (GFP) Pay-what-you-bid is the simplest format of payment and adopted initially by Overture. This mechanism is unstable due to the *Winner's Curse* [21] which causes bids to be changed very frequently as bidders do not want to pay a penny more than what is enough to win the desired position.

Vickrey-Clarke-Groves (VCG) VCG is famous for the property that the dominant strategy of bidders is to bid their true valuation of an item. In VCG keyword auction, the

advertiser i^1 would be charged the externality that he imposes on other advertisers: the difference between the aggregate value all other advertisers would receive if i is not present and the aggregate value all other advertisers would receive if i is present. Despite VCG reduces the incentive for strategising and makes life easier for advertisers, search engines still use GSP today. According to Edelman *et al.* [50], there are two major reasons why it is not adopted by major search engines. First, VCG is hard to understand and advertisers will not stop shading their bids immediately before they realise shading does not work. Second, VCG revenue is at most as large as GSP if all advertisers bid the same amounts under two mechanisms which means search engines will not have incentives to change.

The Laddered Auction In laddered auction, the price for a merchant builds on the bid of each merchant ranked below it [3]. The designer assumes i^{th} merchant also has the i^{th} rank in the auction. For $1 \leq i \leq K$, set the price-per-click p_i charged to merchant i according to the equation:

$$p_i = \sum_{j=i}^K \left(\frac{CTR_{i,j} - CTR_{i,j+1}}{CTR_{i,i}} \right) \frac{w_{j+1}}{w_i} b_{j+1} \quad (2.1)$$

where $CTR_{i,j}$ is click-through rate of i^{th} merchant at j^{th} position, w is the weight on bid b . Aggarwal *et al.* [3] prove its truthfulness such that it is better than GSP for reducing gaming effect and its revenue equivalence with GSP such that it is better than VCG for a higher potential revenue.

GSP with Hidden Cost An important externality in sponsored search is the *Hidden Cost* that one advertiser with low quality or relevance landing page imposes on other advertisers. A poor ad can discourage the searcher from trying to view any sponsored link at all. As a result advertiser's social welfare will be reduced. In economic theory [154, 9, 153], pricing externalities are encouraged to improve social welfare. Similarly, in keyword auction mechanism design, a search engine should encourage ads that give users a positive experience because it makes users more likely to click on other ads [1]. In Abrams's model, the hidden cost per click is denoted by $h_i = h(q_i)$, where q_i denotes the choice of the landing page for advertiser i and b_i denotes his bid. Ranking is run on $b'_i = b_i - h_i$ instead of b_i . Finally, h_i is added to the price per click of bidder i leaving the ranking unchanged.

¹Advertiser i denotes advertiser rank at position i

2.1.2 Ranking Mechanism

Ranking mechanism determines how sponsored link slots are awarded to each advertiser. Although GSP is the industrial standard, the ranking mechanism in GSP varies from one implementation to another. Moreover, different ranking mechanisms result in different user behaviours and search engine revenue, which is why it is an important topic to be discussed in the design of GSP.

Feng *et al.* [55] propose four alternative mechanisms for choosing the allocation of paid slots to advertisers.

- “ v ranking” is also called rank-by-bid which is purely based on the advertiser’s willingness to pay or valuation per click v . The highest payer will be ranked at the top. It is a stylised version of Yahoo’s model before 2009.
- “ $v \times \alpha$ ranking” is based on stylised Google’s model. α represents the expected click-through rate (relevance) of the listing. It is also explained as advertiser effect e_s in [97]. $v \times \alpha$ can be deemed as expected revenue of the listing. The higher the expect revenue, the higher the rank. Therefore it is also called rank-by-revenue.
- “ α ranking” selects the highest k bids and ranks the bidders by their expected click-through rates.
- “Posted Price Ranking” is a mechanism that the search engine sets a reserve price for each position and allocates k highest bidders to k positions and lets them pay the reserved price for that position if their bids meet the corresponding reserves.

“ $v \times \alpha$ ranking” substantially ranks by weighted bids rather than bids themselves. The weight w associated with the bid is usually the advertisement’s expected click-through rate e_s which is determined by a number of factors including ad relevance, historical click-through rate, landing page quality, and etcetera. Pennock *et al.* [97] generalise this model by introducing a family of weights $w_s = e_s^q$ for $q \in (+\infty, -\infty)$. When $q = 0$, advertiser effect is in fact ignored and it becomes Yahoo’s rank-by-bid. When $q = 1$, it covers Google’s rank-by-revenue. They claim that tuning q can significantly improve equilibrium revenue. However, the improvement of revenue may come at the price of future revenue because advertisers and users may be lost due to decrement of their satisfaction.

Lahaie and Pennock [97] extensively investigate the tuning of q based on the revenue

equation at the equilibrium:

$$R(q) = \sum_{s=1}^K \sum_{t=s}^K \left(\frac{e_{t+1}}{e_s}\right)^q e_s (x_t - x_{t+1}) v_{t+1} \quad (2.2)$$

They show that although $q = 1$ yields the efficient allocation, setting of q considerably less than 1 can yield better revenue if v and α correlation is strongly positive. If the bidders are ranked in decreasing order of relevance, then $\frac{e_{i+1}}{e_i} < 1$ and decreasing q slightly without affecting the allocation will increase revenue. Similarly, if bidders are ranked in increasing order of relevance, increasing q slightly will yield an improvement. For perfect positive correlation between value and relevance, $q = 0$ yield more revenue than $q = 1$.

Correlation Between v and α

Besides [97], [55] also identifies that the correlation between bidder values and click-through rates (or relevance) should be a key parameter affecting the revenue efficiency of various ranking mechanisms. However, because their assumptions are different, their findings are different too.

In [55], they have the following findings regarding not just the v and α correlation but the number of competitors N and attention decay factor δ ,

1. Revenue is increasing in the correlation between α and v . The effect is more pronounced for rank-by-bid.
2. Rank-by-revenue strongly dominates rank-by-bid in the region of negative correlation in terms of revenue and weakly dominates it in the region of positive correlation.
3. Both revenue decrease as δ increases.
4. When v and α correlation is strongly positive, the increase of N significantly increases search engine's revenue for both mechanisms. When v and α correlation is strongly negative, N still increases with revenue for rank-by-revenue but decreases with rank-by-bid.
5. When there is a highly positive correlation between α and v , the search engine's expected revenue is approximately concave in the number of sponsored ad links it enrolls. 3 to 7 seems to be the optimal choice. Intuitively, more ad slots will increase the revenue, but too many can negatively affect the overall quality of the search engine and in turn reduce total traffic at the search engine and the clicks on the ads, thereby lowering revenue from sponsored search.

6. The relevance issue does need to be addressed in rank-by-bid. Otherwise, sponsored slots could be awarded to poor relevance ad thus yielding low revenue. In practice, editorial control is conducted to screen out listings below a certain threshold. With a good choice of the threshold value, the modified v ranking outperforms $v \times \alpha$ in the regions of both positive and negative correlation.

2.1.3 Modelling User Behaviours

Two essential factors in click models are position effect and relevance effect. Position effect is how the probability of click depends on position. Relevance effect is how the probability of click depends on the relevance of the listing content. Different arrangements on these two factors generate different click models.

Craswell *et al.* [37] propose four simple hypotheses to model click behaviour over natural search results which also suit sponsored search results. They number position $i \in 1, \dots, N$, denote r_d as the listing's relevance:

$$r_d = p(\text{Click} = \text{ture} | \text{Document} = d) \quad (2.3)$$

denote c_{di} as observed probability of click:

$$c_{di} = p(\text{Click} = \text{true} | \text{Document} = d, \text{Rank} = i) \quad (2.4)$$

Baseline Hypothesis There is no position bias, the probability of clicking on a document at position i is the same as the probability of clicking it at position j .

$$c_{di} = r_d = c_{dj} \quad (2.5)$$

Mixture Hypothesis There is a probability b_i that the user will click on early ranks blindly. The proportion of users who click blindly can be explained by a mixture parameter λ :

$$c_{di} = \lambda r_d + (1 - \lambda) b_i \quad (2.6)$$

Examination Hypothesis This is also called separable click-through rate model. To be clicked, the document link must be both noticed with probability x_i and relevant:

$$c_{di} = r_d x_i \quad (2.7)$$

where x_i is the probability that the document is noticed.

Cascade Model In this model, they assume the user search results from top to bottom. The probability that he/she clicks on one document is dependent on joint probability all document with higher position are skipped.

$$c_{di} = r_d \prod_{j=1}^{i-1} (1 - r_{doc.in.rank:j}) \quad (2.8)$$

They conclude that cascade model fits data the best. But whether it is the best explanation for position effect on sponsored links still remains as a question. Although the separable click-through rate model is currently the popular one and many theoretical works (*e.g.* [50, 171, 98]) are based on it. Kempe *et al.* [87] say it is mainly because of its simplicity for analysis and one of the main drawbacks of separable model is that it completely ignores externalities between ads. However, both anecdotal evidence and user studies [83, 37] suggest that externalities are common.

Another user behaviour to model is the conversion, which is harder because it depends on the complex interaction between the customer and the merchant once he or she is referred to the merchant website by the search engine. Das *et al.* [42] propose an extension of separability model in which the searcher will not only click on ad with high relevance but also convert on ad with higher relevance. Moreover, conversion will come at most from one of the advertisers and the probability of conversion on that ad in slot i is:

$$Pr_{con}(i) = \frac{\theta_i r_i \cdot r_i}{\sum_{1 \leq j \leq k} r_j} \quad (2.9)$$

where θ denotes position effect and r_i denotes the relevance effect. Lately, Patrick Jordan *et al.* propose a conversion model for the use of TAC AA 2009. In this model, more specific keyword phrases get higher chance of conversion and conversion probability is diminished if the advertiser's stock is about to run out.

2.1.4 Analysis of Empirical Data

The handful of empirical studies that exist in search engine marketing have typically analysed publicly available data from search engines. Telang *et al.* [164] summarise that search user behaviours (click and conversion probability) are also affected by the characteristics of keyword itself apart from ad-related characteristics like relevance and rank. They also show that a model that accounts for periodicities fits the search engine visits well. Moreover, different goals of search users can be indicated by search queries. According

to Broder [120] and Jansen *et al.* [78], search goals can be classified into three categories: navigational (*e.g.*, a search query consisting of a specific firm or retailer), transactional (*e.g.*, a search query consisting of a specific product), or informational (*e.g.*, a search query consisting of longer words). Moreover, the length of keyword phrase is also an important determinant of search and purchase behaviour. But anecdotal evidence on this varies across trade press reports. Some studies have shown that the percentage of searchers who use a combination of keywords is 1.6 times the percentage of those who use single keyword queries [90]. In contrast, another study found that single keywords have on average the highest number of unique visitors [77].

Ghose and Yang [61] empirically investigate sponsored search advertisement campaign data from a large retailer in USA. They find that the monetary value of a click is not uniform across all positions because conversion rates are the highest at the top and decrease with rank as one goes down the search engine results page. Though search engines take into account the current period's bid as well as prior click-through rates before deciding the final rank of an advertisement in the current period, the current bid has a larger effect than prior click-through rates. Their analysis shows that keywords that have more prominent positions on the search engine results page, and thus experience higher click-through or conversion rates, are not necessarily the most profitable ones - profits are often higher at the middle positions than at the top or the bottom ones. Besides providing managerial insights into search engine advertising, these results shed light on some key assumptions made in the theoretical modelling literature in sponsored search. They also find that an increase in landing page quality scores is associated with an increase in conversion rates and a decrease in advertiser's cost per click.

Animesh *et al.* [1] look at the presence of quality uncertainty and adverse selection in paid search advertising on search engines. Goldfarb and Tucker [64] examine the factors that drive variation in prices for advertising legal services on Google. Rutz and Bucklin [152] show that there are spillovers between search advertising on branded and generic keywords, as some customers may start with a generic search to gather information, but later use a branded search to complete their transactions. Agarwal *et al.* [2] provide quantitative insights into the profitability of advertisements associated with differences in keyword position and show that profits may not be monotonic with rank.

2.2 GSP Strategy

What advertisers want is the best bidding strategy, which is also the goal of TAC Ad Auction competition. Good return-on-investment ratio is most advertisers' main target regarding the bidding strategy. But there are many other aspects that some particular advertisers care about such as beating the competitors, keeping a certain position on the sponsored link list, sticking with a restrictive budget limit and so on. From game theoretic perspective [137], a good bidding strategy is a stable one with which bid prices converge to an equilibrium where nobody has incentive to change his/her bid any more and the social welfare is maximised at this equilibrium. The remaining of this section will try to cover most well-known bidding strategies that have been proposed, investigated or actually used in the real sponsored search so far.

2.2.1 Greedy Biding Strategies

Greedy bidding strategy is a category of bidding strategies in sponsored search that features maximising the utility return of an advertiser given the expected bids of other advertisers. Assuming that bids are known and other bidders will bid in the next round exactly what they bid in the current round, a greedy bidding strategy [24] for advertiser A is to choose a bid for the next round to maximise his utility relative to this postulated set of bids by other bidders. Because the GSP mechanism allows a range of bid values that will result in the same outcome from A's perspective, A so as to win slot s can choose any bid in the interval (P_s, P_{s-1}) . Therefore greedy bidding strategies have three subcategories where c_s is click-through rate of position s :

Balanced Bidding (BB) in which A's choice of next bid b makes it indifferent between successfully winning the targeted slot s at price p_s , or winning slot $s - 1$ at price b . Formally, A chooses b so that

$$c_s(v - p_s^*) = c_{s-1}(v - b)$$

where v is private value of the advertiser. If s^* is top slot, $b = (v + p_1)/2$.

Competitor Busting (CB) [123] in which A chooses the highest bid value consistent with the target. By playing CB, players focus on forcing their competitors to pay more without affecting their own payment.[192] Formally, A chooses b so that

$$b = \min(v, p_{s-1} - \epsilon)$$

Altruistic Bidding (AB) in which A chooses the lowest bid value consistent with the target slot. Formally, A chooses b such that

$$b = \min(v, p_s + \epsilon)$$

BB is a particular interesting strategy as it has a unique fixed point with payments identical to those of the VCG mechanism. The main convergence properties of BB are the following:

- For two slots, BB converges to its unique fixed point with probability 1.
- For three or more slots, BB does not always converge assuming all players simultaneously update their bids according to BB.
- In an asynchronous model, if and only if one randomly chosen player updates his bid each round according to BB, BB bidding always converge eventually.

Vorobeychik's analysis on bidding dynamics suggests that AB is highly unstable although it yields the greatest pay-offs to players (but least to the auctioneer) [175]. BB has considerable strategic stability but yields lowest pay-offs to players (likely highest search engine pay-offs). Zhou *et al.* [192] is more interested in vindictive bidding (CB) which is also referred as anti-social bidding by [16] because it is likely the most used strategy in the real world. They prove that Pure Strategy Nash Equilibria (PSNE) may not exist when there are at least three players who are all vindictive to each other.

2.2.2 Budget Constrained Strategies

Normally an advertiser's budget is limited. Therefore, they have to solve the problem of placing bids on keywords of their interests so their return can be maximised for a given budget [122]. On both Google and Yahoo, in addition to the bids, the advertiser can also specify a daily budget for each keyword such that every advertiser faces a budget optimization problem. Muthukrishnan *et al.* assume that budget and the set of keyword are already given in single slot case, search companies can predict probability distribution associated with queries in the future with reasonable accuracy and measure the effectiveness of the advertising campaign as the number of clicks. Under these premises, an advertiser has a set T of keywords, with $|T| = n$, and a budget B . For each keyword $i \in T$, they are given $clicks_i$, the number of clicks that correspond to i , and cpc_i , the cost per click of these clicks.

$b_i \in 0, 1$ represents whether or not to bid on keyword i . The optimization objective is to find a solution $\mathbf{b} = (b_1, \dots, b_n)$ with a bid b_i for each $i \in T$ to maximise

$$value(b) = \frac{\sum_{i \in T} b_i clicks_i}{\max(1, \sum_{i \in T} b_i cpc_i clicks_i / B)} \quad (2.10)$$

So maximising $value(b)$ is equivalent to maximising the number of clicks in that the advertiser is under budget, and minimizing the average cost per click the advertiser is over budget. They identify three stochastic versions to model the joint distribution of $clicks_i$ corresponding to different keywords as random variables:

- Proportional Model, in which queries and clicks vary from day to day, but the proportions of clicks for different keywords stay the same.
- Independent Keyword Model, in which each keyword comes with its own probability distribution for the number of clicks, and the samples are drawn from these distributions independently.
- Scenario Model, in which there is an arbitrarily large number of scenarios, each of which specifies the exact number of clicks for each keyword. One scenario is sampled from a given probability distribution over scenarios, determining the numbers of clicks for the problem.

Finally, they show that an optimal fractional solution can be found in polynomial time in the proportional model. For independent model, they present an $2+\epsilon$ approximation algorithm in polynomial time. For scenario model it is NP-hard to reach the optimum.

Zhou and Lukose [191] model the bidding optimization problem as an on-line (multiple-choice) knapsack problem. By achieving a provably optimal competitive ratio, their algorithm can be translated back to fully automatic bidding strategies maximising either profit or revenue for the budget-constrained advertiser. To maximise revenue from sponsored search, their bidding strategy can be ignorant of other bidder's prices and/or click-through-rates for those positions.

Given a daily budget, the advertiser would want to spend the money on the most profitable keywords. There is, however, a trade-off between selecting too few profitable keywords and not spending the entire budget versus selecting too many keywords and depleting the budget too soon, and thus, losing opportunities to receive clicks from more profitable keywords that may arrive later. Rusmevichientong *et al.* [151] formulate a model of keyword selection where they develop an algorithm that adaptively identifies the set of keywords to bid on based on historical performance. The algorithm prioritises keywords based

on a prefix ordering sorting of keywords in a descending order of profit-to-cost ratio. They show that if the cost of each item is sufficiently small compared to the budget, if the expected number of arrivals of any given keyword is not too large, and if the expected number of searched keywords is close to its mean, then the average expected profit generated by the algorithm converges to near-optimal profits. AstonTAC borrows the idea of prefix ordering in the descending order of each keyword's conversion profit.

2.2.3 Heuristic Rules

Heuristic rules are basic strategies that people employ in PPC (Pay-Per-Click) auction. The most popular and fundamental rule is the spend limit over a period of certain length typically one day. In TAC AA [84], a daily spend limit can be set for the whole ad campaign including all keywords or every individual keyword. Several companies market bid management by allowing clients to set higher level rules. Here are some examples from www.gotoast.com:

- *Relative positions*: Allows your listing to always be a certain number of positions above a particular competitor.
- *Gap Jammer*: Moves a bid price to one cent below your competitor. This one is similar to competitor busting (CB) in the family of greedy bidding strategies. However, CB first identifies an optimal position on which the advertiser's own utility is expected to be maximised. If the optimal position is not 1, then the bid is set as high as possible to maximise the competitor's cost.
- *Move to gap*: Move to a gap greater than D in dollars. A 'gap' in an PPC auction exists when there are several groups of bidders whose bid prices are close to each other inside each group and there is a much larger price gap between groups. Those gaps are desirable places to move in because the bidder can avoid the intense competition within a group and subsequent high cost.
- *Time position*: Set a position that you would like to maintain in the auction between certain hours of the day.

Rules are useful for catching exceptions and implementing business logic on branded terms. However, maintaining and adjusting rules for a large set of keywords amidst a continuous, non-stationary auction can be really hard to manage.

2.2.4 Large Scale Optimization

Kitts and Leblanc (2004) [95] develop an optimization-based strategy that also incorporates rules to exert finer control over the behaviour of certain keywords. They test the model through bidding on real search engines. In the same advertising campaign, a human marketing analyst manages the keyword auction in period-1 (01-08-2003 to 14-09-2003) to the best he can and their bidding agent continues to operate period-2 (15-09-2003 to 30-10-2003). In the end, period-2's clicks approximately quadruples period1's with an daily expenditure decrease from \$18.40 to \$16.60.

In their model, they try to find a $b_{k,t}$ to assign to each keyword k , at each future time t . This vector of prices, called the bid allocation plan, should have the property of maximizing the summed profit during the planning horizon, while not exceeding a daily budget. They first estimate unknown functions such as click function conditional on position and position function conditional on bid and unknown parameters such as expected revenue per click r_k using historical data, then set constraints including daily budget, maximum bid and minimum bid as mathematical inequations, finally solve the whole problem as integer programming problem. The estimation of click function needs support of time similarity function to decide weight of historical time t_i 's observation on the forecast of future time t . Instead of simply giving near time a larger weight and far time a smaller weight, they use three kernels to cover hour similarity, day similarity and week similarity. It is interesting because by doing so they automatically assume that users present similar behaviour for the same hour of the day, same day of the week and same week of the year. Consequently, weight on the same hour of different day may be more than the weight on an hour that is further from the hour being forecasted but in the same day.

2.3 CAT Market Design and Analysis

Trading agent competitions are contributive in fostering and exchanging innovative ideas among researchers [181, 159]. After investigating keyword auctions, our research focus is switched to the double auction mechanism by participating in another trading agent competition – the CAT tournament.

2.3.1 Markets Designed by CAT Entrants

An initial representative work on CAT is IAMwildCAT [177] – the winner of the first CAT tournament. It treats the market mechanism design as solving the problem of mul-

multiple trade-offs in the mean time. It introduces the hybrid clearing strategy, which clears shouts to maximise transactions in some rounds and profit in the others. It also takes measures to improve its profit-score conversion ratio. IAMwildCAT also tracks the absolute value of daily overall profit and exploits it when it is small to obtain high profit share score. Moreover, it introduces a side-biased pricing strategy, which gives benefit to the under-represented side depending on the number of buyers and sellers participating in the market. On this aspect, our work is similar but more focused on offering benefits to under-represented intra-marginal traders.

PersianCAT, the winner of CAT-2008, proposes a method for estimating equilibrium price based on market trend [75]. The method can produce more robust estimations than using a sliding window to calculate an average of the recently matched ask-bid pairs [133]. Their work shows the importance of establishing a market on equilibrium price, which leads to positive market outcomes given a sufficiently accurate estimation. In AstonCAT-Plus, we extend PersianCAT's method of equilibrium price estimation [75] by taking into account the latest change in the market on the current day. Besides, Honari *et al.* [75] show that instability in charging fees can cause negative market outcomes when comparing the same markets with almost equal average fees throughout the game. They also show that three factors contribute to the market stability: 1) an equilibrium based pricing policy; 2) an equilibrium based accepting policy; and 3) a stable charging policy.

Lampros *et al.* [158] successfully calculate the global competitive equilibrium as opposite to the equilibrium of any single market. They continually keep track of the highest bids, the lowest asks and the number of goods traded daily, and then form the global cumulative demand and supply curves to compute the desired competitive equilibrium pair of price and quantity. Their method is validated by a mean absolute error of less than 2% to the theoretical global competitive equilibrium. Based on this equilibrium, Mertacor (2008 version) identifies globally intra-marginal traders and implements a strategy to promote the quality of traders in its market. In our AstonCAT-Plus, our method is based on the equilibrium of the single market. From our experiments, it can be seen our approach is not worse than theirs.

Gruman *et al.* use classification techniques to determine the distribution of bidding strategies used by all traders subscribed to a particular specialist. Their experiments show that Hidden Markov Model classification yields the best results. Then, distribution of strategies is used to determine the optimal action in any given game state. Data shows that the GD [62] and ZIP [35] bidding strategies are more volatile than the RE [53] and ZIC [63] strategies, although no traders switch specialists too easily. Finally, they propose

an Markov Decision Process (MDP) framework for determining optimal actions given an accurate distribution of bidding strategies.

Pardoe *et al.* [139] suggest a novel approach to mechanism design. Instead of relying on analytical methods that depend on specific assumptions about bidders, they create self-adapting auction mechanism in which parameters are adjusted in response to past auction results. Their experiments show that the difference between the results of the adaptive method and each fixed choice is statistically significant with 95% confidence according to a paired t-test. What's more, they demonstrate the efficacy of this approach in a situation where a seller must choose from a space of sequential auctions in order to maximise its revenue.

There are other investigations that concentrate on other aspects of market design in CAT. Phelps *et al.* [145] use genetic programming to price transaction at an optimal point in bid-ask spread based on which allocative efficiency can be maximised. Cliff [34] explores a continuous space of auction mechanisms defined by a parametrised version of the continuous double auction, where the parameter represents the probability that a seller will make an offer during any time slice.

2.3.2 Analysis on CAT Specialists

CAT competition is a very important approach to test different designs of double auction mechanism. It alone usually cannot provide a complete view of the relative strengths and weaknesses of different specialists because the performance of a specialist in the competition depends upon the composition of its opponents [132]. Hence, it is necessary to carry out post-tournament analysis. [129, 130] provide a refined classification of the CAT-2007 entries based on their internal designs to the taxonomy covered by [187]. Using white-box analysis, they attempt to relate market dynamics to the auction rules adopted by these entries and their adaptive strategies through a set of post-tournament experiments. Using black-box analysis, they reveal the strength and weakness of the specialists in several scenarios and demonstrated some vulnerabilities in entries placed highly in the competition.

Besides, Cai *et al.* [23] examine how standard economic measures, like allocative efficiency, are affected by the presence of multiple markets for the same goods. They find that dividing traders between several small markets typically leads to a lower efficiency than grouping them into one large market. Nevertheless, the movement of traders between markets and price incentives for changing markets can reduce this loss of efficiency. [134] also finds that generally the traders are attracted towards lower charging markets and these

markets generate more profit.

In [129], the CAT developers introduce a specialist names MetroCat, which is developed based on the insights about the CAT game: 1) it is crucial to maintain a high transaction success rate; 2) registration and information fees should be avoided; 3) intra-marginal traders still stay with the market as long as they make a considerable amount of profit through transactions after covering fees. Specifically, they use a history-based shout accepting policy developed based on GD trading strategy [62], which estimates the probability a shout would be matched, and only accepts those shouts with a probability higher than a specified threshold. They demonstrate such a market can successfully beat all entries in CAT 2007. However, their clearing strategy is not adaptive. Rather, our specialist AstonCAT-Plus can adapt its clearing thresholds according to the change of a market status. Although we did not run experiments to compare our market agent with those in CAT 2007, generally speaking adaptive one should perform better than non-adaptive ones. Moreover, Niu *et al.* [129] demonstrate some vulnerabilities of entries that placed highly in the 2007 CAT competition and provide a general approach to conducting experimental analysis of similar competitive games.

Additionally, [157] particularly focuses on how market agents use registration fee policies to attract intra-marginal traders and drive out extra-marginal traders. They also study how Nash equilibrium changes across two markets when one of them charges a registration fee. Robinson *et al.* [149] make an effort in analysing the generalisation abilities of specialists in CAT 2008 and show that specialists are sensitive to several factors in the competition, including the trading strategies distribution and scoring period. Compared with the above investigations, our post-tournament experiments and analysis emphasise on revealing the essential features of success specialists, such as the balanced trader structure and the market share dominance regardless of trader mixtures.

2.4 Agent-based Financial Market Simulation

Our investigation of keyword auction and double auction is largely based in multi-agent simulation systems. With these experience, we build a simulation of financial market using an alteration of double auction model which particularly reflects the inconstant behaviour of financial traders in terms of changing trading directions. Here we discuss the state of the art in the field of computational economics.

2.4.1 General Computational Finance

Agent-based computational economics (ACE) [167, 101], alternatively the micro simulation [103] is an important research field for understanding complex patterns, stylized phenomena observed in economic systems especially financial trading systems. ACE emphasises the need to represent traders as individual agents in order to study the way macro features emerge from individual interactions [15, 103]. It is able to tackle some limitations of the analytical models in economics and finance [117]. Agent-based approaches attempt to model the market as evolving systems of competing, autonomous interacting agents and emphasise their learning dynamics [168, 99]. ACE has been successfully applied to various economic studies [94, 6, 8, 4] and the most popular area should be computational finance [170]. In this background, we are strongly motivated to create a financial market simulation based on our experience in the research of double auction market and a specific angle of viewing the financial market.

Over the last 30 years, there is growing literature attempting to model financial interactions from the agent perspective in which the most influential work is probably the Santa Fe Artificial Stock Market Institution (SFI) [7]. The SFI framework [101] employs a rational expectations asset-pricing model, genetic algorithm learning, and Walrasian tatonnement as the market clearing mechanism. Despite its influence, Santa Fe market is not the first and has quite some contemporaries [36, 91, 58, 46]. These are followed by much of the later literature in being concerned with the interaction of potentially destabilizing trend following strategies and their interactions with others. [104] is another market example which introduced constant relative risk aversion preferences along with varying agent memory lengths. Moreover, the market of [11] introduce trading strategies based on neural networks. [110, 93] and [31] provide relatively simple tractable dynamic frameworks for agent interactions.

In the last ten years or so, more agent-based models have been proposed. Some focus on agent strategies or characteristics, *e.g.*, [160, 116, 51, 72]. For example, Yang [188] investigates the convergence property in a double auction market where artificial neural networks take on the role of traders, who form their expectations about the future return and place orders based on their expectations. The author finds that the convergence of market price is sensitive to the deviation from rationality of agents. This is quite similar to one of our findings with BDA market: market efficiency largely comes from rational selection of trading directions. He *et al.* [72] develop an algorithm that employs heuristic fuzzy rules and fuzzy reasoning mechanisms in order to determine the best bid to make given the state

of a continuous double auction market. Moreover, they show how an agent can dynamically adjust its bidding behaviour to respond effectively to changes in the supply and demand. The adaptive-aggressiveness strategy [176] is based on a short-term and a long-term learning of the agent's bidding behaviour. For the short-term learning, the motivation was to immediately respond to fluctuations in the market conditions. For long-term learning, the motivation is to respond to more systematic changes in the market conditions *i.e.*, market shocks.

Some focus on market mechanisms, *e.g.*, [147, 178, 29, 132, 139]. For example, Walia *et al.* [178] use a self-adaptive Evolutionary Strategy (ES) to explore the space of possible auction types in a CDA populated by Zero Intelligence Constrained traders and demonstrate that non-standard variants of the CDA can provide favourable dynamics for trading strategies. Some borrow ideas from statistical mechanics, *e.g.*, [105, 111]. Some modelling frameworks are inspired by Minority Games, *e.g.*, [27, 156]. For example, Challet *et al.* use the minority game model to study a broad spectrum of problems of market mechanism. The central issue they are concerned about is the information flow: producers feed in the information whereas speculators make it away. They claim market impact is shown to play an important role and a strategy should be judged when it is actually used in play for its quality.

In terms of whether portfolio management is involved in modelling artificial financial market, we further detect two groups of paper: (i) mono-asset market [88, 86, 56, 127] and (ii) multi-asset market [109, 89, 32, 25]. In addition, Brock and Hommes [20, 74] model financial market as an adaptive belief system and point out a large fraction of chartists tends to destabilize the market. Lux and Marchesi [112, 113] use mass-statistical approach to represent traders and model agents in groups, switching the proportion of agents over the alternatives in a stochastic manner. Danials *et al.* [41] describe a microscopic dynamical statistical model for the continuous two-sided auction under the assumption of IID random order flow. In the model of Franci *et al.* [57], traders do not only analyse historical data but also take into account the forecast of other successful agents and news in order to formulate their own trading intentions.

Apart from what's mentioned above, there are recent work regarding financial market simulation. Raberto *et al.* [147] create Genoa market in which they investigate stylised facts of the limit order book and the distribution of waiting times between two consecutive transactions. LeBaron [100] introduces a new framework extended from the pioneering SFI mentioned above, which emphasises memory length as a key heterogeneity dimension, and determining factors for convergence to homogeneous equilibrium pricing instead of

learning speeds. Andre *et al.* [26] develop AgEx market to address the demand of trading strategies' development and assessment in financial market. Veryzhenko *et al.* [172] present ArTificial Open Market API (ATOM), which is developed as a large scale experimental platform with generic architecture and heterogeneous agents populations. They demonstrate a series of key points and principles that have governed the development of an agent-based financial market in the form of an API.

In terms of modelling financial market, literature in quantitative finance has to be mentioned.

The prevalent theory of financial markets during the second half of the 20th century has been the efficient market hypothesis (EMH) which states that all public information is incorporated into asset prices. Any deviation from this true price is quickly exploited by informed traders who attempt to optimize their returns and it restores the true equilibrium price. For all practical purposes, then, market prices behave as though all traders were pursuing their self-interest with complete information and rationality.

2.4.2 Double Auction Financial Market

When designing an agent-based market system, an important issue is how the price of the good is formed. It can be given exogenously or generated endogenously by the interaction of autonomous trading agents. Automated electronic markets with endogenous prices frequently use auction protocols as the mechanism to determine prices of the goods to be traded [141]. Double auction protocol is normally adopted to build both real financial market and its simulations. Our BDA market project is an effort to simulate financial market using a double auction model, in which we also create a high-performance kernel trading strategy. Hence, our research is related to two types of literature: agent-based computational simulations [183] which emphasises the model's ability to replicate the real market and double auction mechanism and strategy design which emphasises the model's ability to allocate resource efficiently.

From the aspect of double auction mechanism and strategy design, there are quite a few previous literature that paves the way for our research of BDA market. Gode and Sunder [63] claim that allocative efficiency of a double auction derives largely from its structure so that "zero-intelligent" traders imposing budget constraints are sufficient to raise the allocative efficiency of these auctions close to 100%. Cliff [33] claims that more than zero-intelligence is required to achieve efficiency close to that of markets with human traders. [70, 115] study the effects of changes in supply and demand in the context of the New

York Stock Exchange and [62] considers the performance of artificial trading agents under varying conditions. Recently, TAC CAT tournament [22] opens a door to the exploration of the optimal design of double auction market using competition approaches. However, Chang *et al.* [30] point out that the success of a particular design of double auction market is not independent of the distribution of trading strategies equipped by the traders. They also discuss how to make a double auction market more efficient, robust and profitable by setting effective rules from various aspects in the background of CAT tournament [29].

Additionally, the extended Glosten and Milgrom microstructure model [44] focuses on the market makers' quote-adjusting strategy. In this model, non-parametric density estimation technique is proposed for maintaining a probability distribution over the expectation of real stock value that market-maker can use to set prices. Our *Kernel* trading strategy employs the same technique to maintain two probability distributions over the expected transaction of both ask and bid. Blum *et al.*, [12] examine the design of matching algorithms with good worst-case performance within the framework of competitive analysis. Bredin *et al.* [17] construct a general framework that facilitates a truthful dynamic double auction from truthful, static double auction rules. Zhao *et al.* [190] develop computational efficient matching algorithms using weighted bipartite matching in graph theory. Muchnik *et al.* [121] introduce continuous time asynchronous model to simulate financial market under the name of NatLab (asynchronous double auction). Simulations show that market dynamics can be drastically changed by a small fraction of trend followers. Our BDA market also features asynchronous order submission but our timing discrete currently.

In terms of trading strategy design in double auction market, we have reviewed many popular existing ones including zero intelligence [63], ZIP [35], GD [62], RE [53], GDX [165], Fuzzy logic based strategy [72], Adaptive-Aggressive bidding strategy [176] and so on. Their details are given in Section 5.1.5. In above literature, trading directions of individual agents in a double auction market are either generated randomly or pre-defined, which drives us to create a market with dynamically-decided rational trading directions of the agents.

In terms of making decisions in real financial markets, there are even more strategies developed based on the data of real financial markets or high-level realistic simulation of stock or foreign exchange markets such as SFI. Most of them use machine learning algorithms combined with technical analysis indicators [106, 13, 189, 124] to generate trading decisions. Neural network is a commonly studied approach [169, 92, 182]. Even though these studies indicate that they outperform their benchmarks, the main problem with the neural network approach is the difficulty of interpreting the trading rules generated, espe-

cially in the case of complex networks with many nodes and hidden layers. In comparison, genetic algorithm has the advantage that the rules are interpretable. For example, Lettau [102] builds an agent-based financial market using simple agent benchmarks based on genetic algorithms. Routledge [150] extends the basic framework of Grossman and Stiglitz [68] with agents that can learn by using genetic algorithm. Reinforcement learning is another popular approach. For example, in [10], order flow data is coupled with order book derived indicators and pattern recognition techniques are employed to infer trading strategies on the underlying time series. They show that using order flow and order book data is usually superior to trading on technical signals alone. Moriyama et al. [119] successfully test the application of reinforcement learning to trade on a futures market simulator (U-Mart) of the large Japanese industrial companies (J30) index. In addition, support vector machine is applied by Tay and Cao [162] giving more weight to more recent values and show improved forecasting results of the S&P 500 index, and US and German government bond futures than using moving averages and lagged prices.

In contrast to that many strategies take input of low frequency financial data *e.g.*, daily prices and returns, Creamer [38] proposes a high frequency trading strategy for equity index futures. His agent uses the expert weighting algorithm [39] to forecast a price trend as an input for a trading strategy based on a variation of a market maker strategy proposed by [40]. It automatically calibrate a trading model using different versions of the same technical analysis indicators because their approach takes advantage of boosting's feature selection capability to select an optimal combination of technical indicators and generates experts at different moments of the trading cycle. The main objective of their algorithm is to predict the return of the future contract in the next period. It combines the output of several experts and suggests a short or long position. If the expected position is positive (negative), the trading agent sends a buy (sell) limit order at prices slightly lower (higher) than the bid price at the top of the buy (sell) order book less (plus) transaction costs. In comparison to our BDA market, their trading directions signals more realistic because they do rely on an assumed private value which still needs to be derived from other information in practice. However, the general trading action structure is similar to ours, which is to generate a trading direction first, then a price of order.

Chapter 3

Adaptive Strategy Design for TAC Ad Auction

Sponsored search is one of the most cost-effective and efficient ways of advertising because advertisers only pay when a user shows true interest in their advertised product by clicking on their ads, which makes it the most popular advertising manner of search engines. The success of such advertising approach should attribute to the introduction of GSP which fairly prices ad slots through competition and efficiently determines allocations of advertisement resources. This chapter discusses our research into the strategy design problem in GSP. Specifically, we will describe TAC AA competition and explain how AstonTAC – one of the most successful strategies in TAC AA is designed, implemented and evaluated. In order to help readers understand this chapter and possibly replicate our agent, we provide Appendix A to show our agent and experimental settings and Appendix B to list symbols used in our agent design and description.

3.1 Sponsored Search and Ad Auction Competition

The investigation of sponsored search generally falls into three categories: 1) search user behaviour modelling; 2) mechanism design faced by the search engine; 3) strategy formulation faced by the advertiser. The Trading Agent Competition Ad Auction¹ provides an ideal test bed for advertiser strategies. In TAC AA, there are three kinds of agents: users, publishers and advertisers. The behaviour of the users and publishers are generated by TAC AA server according to some fixed stochastic policy [84]. There are eight advertiser agents

¹<http://aa.tradingagents.org>

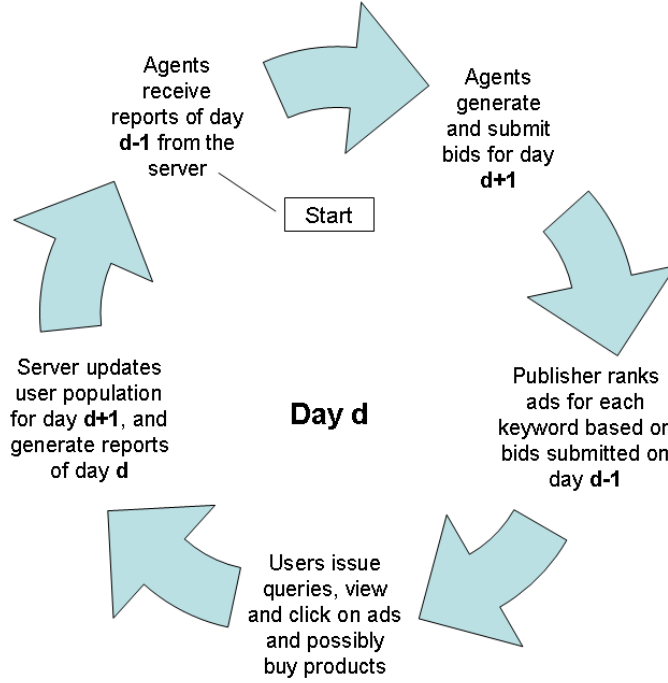


Figure 3.1: Daily activities cycle in TAC AA.

(entrants to the competition) that compete against each other on ad placement for search result given queries over 60 days of bidding periods.

In more details, agents represent retailers of home entertainment products featured by three manufacturers (*Flat*, *Lioneer* and *PG*) and three components (*TV*, *Audio* and *DVD*). Altogether, there are nine distinct products. A query generated by a user is a (*manufacturer*, *component*) pair and unspecified manufacturer or component is denoted as '*null*'. In total, there are 16 possible queries at three focus levels denoted by F_0 , F_1 , F_2 . The more specific the query, the higher the focus level.

$$Q_{F_0} : \{(null, null)\}$$

$$Q_{F_1} : \{(Flat, null), (Lioneer, null), (PG, null), (null, TV), (null, Audio), (null, DVD)\}$$

$$Q_{F_2} : \{(Flat, TV), (Flat, Audio), (Flat, DVD), (Lioneer, TV), (Lioneer, Audio), (Lioneer, DVD), (PG, TV), (PG, Audio), (PG, DVD)\}$$

On each day of the game and for each query type above, an auction is run to determine the ad placements. Once an ad is clicked and leads to a customer transaction, it is called

a conversion. Based on the focus levels, distributions of click probability P_{click} , continuation probability $P_{continuation}$ (probability that a user proceeds to click the next ad), and conversion rate $P_{conversion}$ differ between queries. In addition, each advertiser is assigned a Manufacturer Speciality (MS) and a Component Speciality (CS) in each game. If a query matches MS , it will receive a high conversion value. If a query matches CS , it will receive a high conversion rate. The number of conversion is softly constrained by *Distribution Capacity* C^{cap} . TAC AA introduces C^{cap} to impose the effect of diminishing marginal value [84] of conversion: when the number of conversion exceeds C^{cap} , $P_{conversion}$ starts to drop and result in lower conversion profit due to the increased cost.

On each day of the game and for each query type, the advertiser agent submits a bid to the publisher. Such a bid specifies the bid price (the maximum amount that an advertiser is willing to pay for a click on his ad), the spend limit (the corresponding ad will be excluded once the spend limit is reached) and the ad display type (either *generic* or *targeted*). At the end of the game, agents are evaluated based on their cumulative surplus: sales profit less cost they paid for all the clicks received in the game.

Figure 3.1 illustrates activities on each day of the game. On day d , each advertiser agent first receives market reports of day $d-1$, then decides a bid for each query to submit for day $d+1$. The publisher ranks bids submitted by different advertiser agents on day $d-1$ and works out a price per click to charge on each query for each agent. As users click on ads and buy products from advertiser agents, the server collects every agent's impression, position, clicks, conversions, revenue and cost to generate market reports of day d . Bidding period begins from day 0 and agents' first bid submission is for day 1. Consequently, market reports are available from day 2. Therefore, the first market report is for day 1 and accessible from day 2.

According to specific settings of TAC AA, designing a successful agent mainly faces two challenges: **a)** without knowledge of bid prices of other participants, what is the optimal position and how to decide an appropriate bid price to target it? **b)** given indeterministic conversion limit, how to maximise profit in terms of both number of conversion and conversion value? To address the first problem, we find an alternative way to build bid prices on Market-based Value Per Click (MVPC) which can be estimated based on system parameters or market reports. To overcome the second challenge, we estimate the true maximum number of conversion allowed by distribution capacity and select only the most profitable queries to bid on and fill the expected conversion allowance every day.

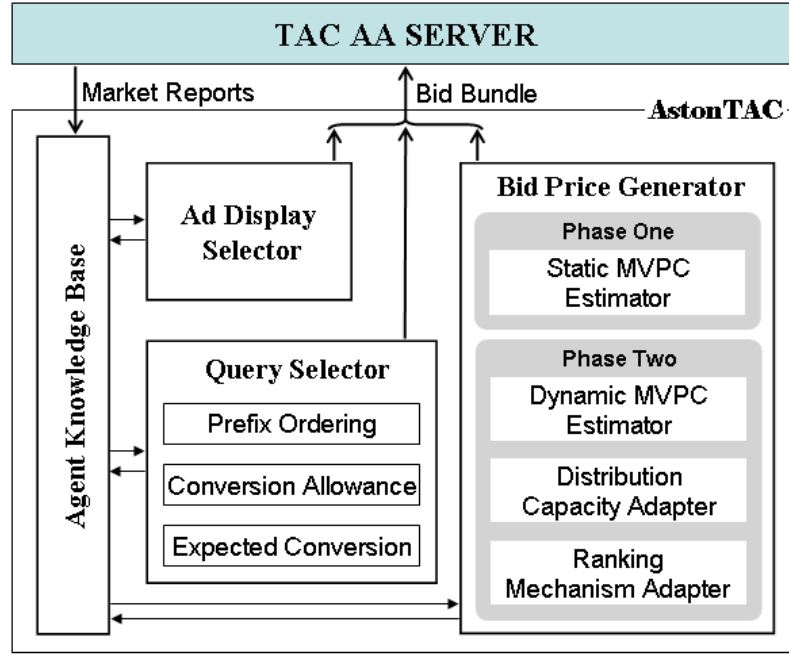


Figure 3.2: Architecture of AstonTAC.

3.2 AstonTAC

In this section, we look into details of the design and evaluation of AstonTAC. AstonTAC is the runner-up in the Ad Auction Game of 2009 International Trading Agent Competition. In particular, we focus on how AstonTAC generates adaptive bid prices according to the Market-based Value Per Click and how it selects a set of keyword queries to bid on to maximise the expected profit under limited conversion capacity. Through evaluation experiments, we show that AstonTAC performs well and stably not only in the competition but also across a broad range of environments.

AstonTAC is composed of four components: Agent Knowledge Base, Bid Price Generator, Query Selector and Ad Display Selector shown in Figure 3.2.

Agent Knowledge Base is designed to process, organise and record information from the server and turn them into knowledge for other components to use. The information it deals with can be divided into two categories: static information including setting parameters and game initialisation information and dynamic run-time information such as market reports received on daily basis. The other three components collectively generate the bid $B = \langle b, l, t \rangle$ for each query q and submitted as a bundle on each day of the game. In particular, Bid Price Generator (Section 3.2.1) calculates a bid price b . Query Selector (Section

3.2.2) specifies a spend limit l . Ad Display Selector (Section 3.2.3) makes the choice of ad display type t .

3.2.1 Bid Price Generator

Based on the availability of market reports, the whole game can be divided into two phases: Phase One (day 0 and 1) and Phase Two (from day 2 to 59). Bid price generation in both phases is based on the same concept - Market-based Value Per Click (MVPC) - of each query.

Definition 1. A query's MVPC is the expected conversion revenue minus advertising cost that a click on its ad incurs.

We introduce MVPC because we find clicks a kind of special commodity as they are not keepable. In an auction of normal commodity, the bidder is willing to pay up to his/her valuation for an item because he/she can own the item after the payment. However, in a keyword auction, clicks cannot be owned so that bidders always want to make a profit out of them. For this reason, we believe MVPC is the true worth of a click to an advertiser. Unlike the conventional Value Per Click [50, 174] which is the expected conversion profit of a click, MVPC is much more dynamic by incorporating the advertising cost. Basically, if it is assumed that revenue-per-click is independent of position [95] which is disapproved by [61], MVPC still varies as cost-per-click (CPC) is expected to vary for different positions. Moreover, due to conversion limit, one query's conversion affects the others' conversion rates rendering the change of revenue-per-click such that MVPC reflects the real value of a query in an interdependent multi-query environment. Therefore, if we can estimate every query's MVPC and set a bid price accordingly, our bids should automatically approximate the best response to the market without explicitly targeting any position. MVPC is estimated in two different ways for different phases of the game.

Phase One

In phase one, MVPC is estimated using expected revenue-per-click multiplied by a fixed discount ratio $r_{discount}$ indicating the proportion of profit in revenue. The expected revenue-per-click is a product of expected conversion value and conversion rate. Because MVPC is based on static fundamental information, it is denoted as $v_{static,q}$,

$$v_{static,q} = P_{conversion,q} \cdot v_{con,q} \cdot r_{discount} \quad (3.1)$$

The conversion value $v_{con,q}$ in above formula is given by,

$$v_{con,q} = \begin{cases} USP(1 + MSB) & \text{if } q_m = MS \\ \frac{2}{3}USP + \frac{1}{3}USP(1 + MSB) & \text{if } q_m = null \\ USP & \text{otherwise} \end{cases} \quad (3.2)$$

where USP (Unit Sale Profit) is the standard conversion value, $MSB=50\%$ is MS bonus rate [84] and q_m denotes the manufacturer part of a query. The first line of the above formula means $q_m=MS$ queries receive a 50% higher conversion value than the normal USP . The second branch means $q_m=null$ queries has 1/3 probability to receive a higher conversion value and 2/3 probability to receive a normal conversion value. The third branch means if a query's manufacturer part is a specific manufacturer other than MS , its conversion value will be exactly USP .

$P_{conversion,q}$ is calculated in a similar way,

$$P_{conversion,q} = \begin{cases} \frac{(1+CSB)\pi_q}{1+CSB\pi_q} & \text{if } q_c = CS \\ \frac{2}{3}\pi_q + \frac{1}{3}\frac{(1+CSB)\pi_q}{1+CSB\pi_q} & \text{if } q_c = null \\ \pi_q & \text{otherwise} \end{cases} \quad (3.3)$$

where q_c denotes the component part of a query, $CSB=50\%$ is CS bonus rate and π_q denotes the baseline conversion rate of query q which is determined by q 's focus level. $F0$ queries' baseline conversion rate is 0.1. $F1$ queries' baseline conversion rate is 0.2. $F2$ queries' baseline conversion rate is 0.3. On this basis, if $q_c=CS$, a superior conversion rate calculated by a η function $\eta(x, y) = \frac{xy}{xy+(1-x)}$ where $x=\pi_l$ and $y=(1+CSB)$ will apply. If a query's q_c is $null$, it has 1/3 chance to receive the superior conversion rate because generally 1/3 user population's preference matches the advertiser's CS . If a query's q_c is a specific component other than the advertiser's CS , its conversion rate will be exactly its corresponding π_l .

The discount ratio $r_{discount}$ in Formula 3.1 is a value between 0 and 1. The best $r_{discount}$ is chosen so that AstonTAC's accumulative profit in the first five day of the game is maximised in experimental games.

By now, $v_{static,q}$ has been found. Apparently, queries with relatively high $v_{static,q}$ are expected to produce more profit on each conversion and vice versa. Since slightly higher bid price does not necessarily raise the cost due to the features of Generalised Second Price [50, 54] mechanism, it is sensible to set high bid prices for relatively high-value queries to increase their chance of receiving top positions. To counterbalance the possible increase of conversion from high-value queries under restricted total conversion, relatively low bid prices should be set for relatively low-value queries. To this end, $h_{v,q}$ ($80\% \leq h_v \leq 120\%$) is introduced to enforce the heuristic. In addition, higher C^{cap} means more conversions

and clicks are acceptable. Therefore, we also introduce $h_c \in \{85\%, 100\%, 115\%\}$ so that AstonTAC generally bids lower when $C^{cap} = 300$ or higher when $C^{cap} = 500$. Together bid of each query in phase one is given by,

$$b_{0,q} = v_{static,q} \cdot h_{v,q} \cdot h_c \quad (3.4)$$

where $h_{v,q}$ varies with the static MVPC of q and h_c varies according to the assigned distribution capacity from game to game.

Phase Two

In phase two, dynamic MVPC of query q is denoted by $v_{dynamic,q}$, which is calculated according to the dynamic market reports. Consequently, a query's bid b_q in phase two is built on its $v_{dynamic,q}$ and adjusted by distribution capacity adapter δ and ranking mechanism adapter β_q simultaneously:

$$b_q = v_{dynamic,q} \cdot \delta \cdot \beta_q \quad (3.5)$$

Dynamic MVPC Formula 3.6 incorporates information about a query's revenue, cost and clicks from recent W days, where W is the size of aggregation window for distribution capacity (Specifically, $W = 5$), to calculate the average profit of a click so far as the expected profit a click can possibly make on the next day:

$$v_{dynamic,q} = \frac{\sum_{i=d-1}^{d-W} revenue_{q,i} - \sum_{i=d-1}^{d-W} cost_{q,i}}{\sum_{i=d-1}^{d-W} click_{q,i}} \quad (3.6)$$

We aggregate W -day data for two reasons: (i) C^{cap} takes effect on the basis of W -day aggregate number of conversion. (ii) it reduces the unwanted fluctuation of $v_{dynamic,q}$ caused by system dynamics. As a result of the above formula, $v_{dynamic,q}$ is highly responsive and adaptive to the change of environment represented by three key factors:

1. *Search user population* - deciding the baseline number of possible clicks, conversion and revenue
2. *Competition intensity* - deciding ad rank and cost-per-click
3. *Conversion probability* - deciding number of clicks needed to generate a conversion

The rationale behind is that we map the above three factors to the three components *revenue*, *cost* and *clicks* respectively in Formula 3.6. *Revenue* is an indicator of user population - a larger number of active users tend to generate more revenue assuming other two

factors are fixed. *Cost* is an indicator of competition intensity - assuming user population is fixed and C^{cap} is not filled, the severer the competition, the more it costs to maintain the same position. Finally, *number of clicks* is an indicator of the conversion probability - excess conversion beyond C^{cap} causes lower $P_{conversion}$, which means more clicks are needed to generate the same number of conversions for any query.

Distribution Capacity Adapter Distribution Capacity Adapter is designed to adapt our bid prices to $C^{cap} \in \{300, 400, 500\}$ a decisive factor in TAC AA explicitly. It strongly confines the advertiser's potential profit by affecting conversion rates. According to the focus levels, each agent's each query is set a default $P_{conversion,def}$ by the server. During the game, once W-day accumulative conversion exceeds C^{cap} , the timely conversion rate $P_{conversion,t}$ will start to drop below $P_{conversion,def}$. When $P_{conversion,t}$ is sufficiently low, clicks will make losses rather than profits because users only click on the ad but almost never make a transaction. Hence, there is a trade-off between number of conversions and conversion rates at which a critical number of conversion $C_{crit}(C_{crit} > C^{cap})$ occurred such that profit of next conversion equals to its cost. The aim of setting a bidding constraint is to keep them both high so that accumulative profit can be maximised. C^{cap} is a soft constraint because exceeding C^{cap} does not stop conversion but only reduce its probability. C_{crit} can be estimated (see Section 3.2.2) but cannot be used here because bids are only allowed to be changed daily rather than every time a conversion happens. Eventually, the ratio δ between C^{cap} and expected W-day aggregate conversion by day d denoted by $c_{agg,d}$ is found to be a suitable indicator of bid adjustment over $v_{dynamic}$. To obtain δ , $c_{agg,d}$ is first estimated using weighted average,

$$c_{agg,d} = \frac{\sum_{t=d-1}^0 (w_t \sum_{i=t}^{t-(W-1)} c_i)}{\sum_{t=d-1}^0 w_t} \quad (3.7)$$

where exponential weight $w_t = \omega^{d-t-1}$ ($0 < \omega < 1$) and c_i denotes total conversion from all queries on day i . Then the adjustment factor δ is given by,

$$\delta = \frac{C^{cap}}{c_{agg,d}} \quad (3.8)$$

The intuition behind δ is: if $\delta > 1$, C^{cap} is expected to be under-filled on day d , then all bids are increased by δ , then the number of clicks is expected to increase as well as the subsequent conversions on day $d+1$; if $\delta < 1$, all bids are reduced for the opposite effect. In this game, δ falls in the range between 0.6 and 1.8.

-
1. Prefix-ordering queries.
 2. Estimate conversion allowance $C_{w,d+1}$ for day $d+1$.
 3. Estimate expected conversion of each query $c_{q,d+1}$ on day $d+1$.
 4. Identify a bidding set of queries $A = \{1, \dots, s\}$, $s \in \{1, 16\}$:
 - 4.1. IF $C_{w,d+1} \leq 0$ THEN $A = \phi$; GOTO 5.
ELSE initialise s to 1.
 - 4.2. WHILE $\sum_{q \in A} c_{q,d+1} < C_{w,d+1}$ and $s < 16$
DO $s = s + 1$.
 - 4.3. $A = \{1, 2, \dots, s\}$.
 5. Set $l = \infty$ for each $q \in A$ and $l = 0$ for each $q \notin A$.
-

Table 3.1: The algorithm of query selector.

Ranking Mechanism Adapter Ranking mechanism adapter adjusts the bid price further by taking into account of the ranking mechanism adopted by the publisher. TAC AA employs a squashing parameter χ ($0 \leq \chi \leq 1$) initialised at the beginning of each game to interpolate between two extremes: $\chi = 0$ is equivalent to rank-by-bid and $\chi = 1$ is equivalent to rank-by-revenue [97]. Specifically, given e_q as the estimated click through probability by the publisher for query q and b_q as the bid on q , the ranking score is calculated as $b_q(e_q)^\chi$. In order to adapt our bid prices to the dynamic ranking mechanism ranking mechanism adapter β is introduced and unknown e_q (e_q employed by the publisher is not revealed to the advertiser) can be estimated using the aggregate Click-Through-Rate (CTR) and denoted as e'_q ,

$$\beta_q = (1 + e'_q)^{-\chi} = \left(1 + \frac{\sum_{i=0}^{d-1} Click_q}{\sum_{i=0}^{d-1} Impression_q} \right)^{-\chi} \quad (3.9)$$

Based on β_q , bid price b_q stays unchanged if $e'_q = 0$, or $\chi = 0$. Otherwise, b_q is reduced with increase of either e'_q or χ .

3.2.2 Query Selector

Given limited conversion capacity, query selector selects only a set of queries to bid on such that the expected available conversions are allocated to the queries that can potentially generate high profit. The following table shows the selection process.

As we can see, the output of query selector is a set of spend limits. Selected queries are not restricted by a spend limit so that their spend limit is set to infinite. Unselected queries will not be active on day $d+1$ so that their spend limit is set to zero. The intuition here is:

by estimating conversion allowance, the maximum number of conversion before expected conversion profit drops to zero is revealed; by prefix-ordering and selecting first s queries to put in the bidding set A, we make sure the allowance is used by high-profit queries only. How Step 1, 2 and 3 work is explained in the remaining of this section.

Prefix-ordering Prefix-ordering [151] is used to sort and prefix queries in the descending order of their profit-per-conversion (PPC). A query's PPC is calculated as follows,

$$PPC_q = \frac{\sum_{i=d-1}^0 \text{revenue}_{q,i} - \sum_{i=d-1}^0 \text{cost}_{q,i}}{\sum_{i=d-1}^0 \text{conversion}_{q,i}} \quad (3.10)$$

where values of variables are obtained from market reports.

Conversion Allowance Referring to the discussion of distribution capacity adapter in Section 3.2.1, expected conversion allowance is the difference between C_{crit} and conversions of four recent days including day d which can be estimated as $c_{agg,d} - c_{d-4}$, ($c_{agg,d}$ is given by Formula 3.7). Once C^{cap} is exceeded, every additional conversion lowers timely conversion rate $P_{conversion,t}$ by $\lambda = 0.995$. Hence, we model C_{crit} as $C_{crit} = C^{cap} + n$ where n is the number of additional conversions by which $P_{conversion,t}$ reaches a critical value $P_{conversion,crit}$ such that expected conversion revenue equals to conversion cost (clicks needed to generate a conversion times CPC). At this equilibrium point, n is maximised and additional conversion will start to make a loss rather than profit. Since both conversion revenue and CPC differ across queries, we introduce general conversion revenue v'_{con} as the average revenue with respect to total conversion and general CPC c'_{click} as average cost with respect to total clicks from all queries,

$$v'_{con} = \frac{\sum_{i=d-1}^0 \sum_{q \in all} \text{revenue}_{q,i}}{\sum_{i=d-1}^0 \sum_{q \in all} \text{conversion}_{q,i}} \quad (3.11)$$

$$c'_{click} = \frac{\sum_{i=d-1}^0 \sum_{q \in all} \text{cost}_{q,i}}{\sum_{i=d-1}^0 \sum_{q \in all} \text{click}_{q,i}} \quad (3.12)$$

Mathematically, the equilibrium point can be presented as,

$$v'_{con} = P_{conversion,crit}^{-1} \cdot c'_{click} \quad (3.13)$$

Since $P_{conversion,crit}$ equals to $P_{conversion,std} \times \lambda^n$, we have,

$$n = \log_{\lambda} \frac{P_{conversion,crit}}{P_{conversion,std}} = \log_{\lambda} \frac{c'_{click}}{v'_{con} \cdot P_{conversion,std}} \quad (3.14)$$

where $P_{conversion,std}$ is an average of baseline conversion rates weighted by the distribution of both queries and search population towards different focus levels. n decreases with c'_{click} because the larger the conversion cost the less excess conversion is needed to reach the equilibrium point. n increases with v'_{con} because the larger the conversion revenue, it takes more excess conversion to bring $P_{conversion,t}$ down from $P_{conversion,std}$ to $P_{conversion,crit}$. Once n is found, conversion allowance $C_{w,d+1}$ is available too,

$$C_{w,d+1} = C^{cap} + n - (c_{agg,d} - c_{d-4}) \quad (3.15)$$

Expected Conversion We model expected conversion of each query on day $d+1$ as a product of expected impression, click probability and conversion rate,

$$c_{q,d+1} = impression_{q,d+1} \cdot P_{click,q,d+1} \cdot P_{conversion,q,d+1} \quad (3.16)$$

$impression_{q,d+1}$ is estimated based on impressions occurred on query q in last $Pr_{burst}^{-1} = 10$ days where Pr_{burst} is the search population burst rate. $P_{conversion,q,d+1}$ is estimated as a product of $P_{conversion,q}$ given by Formula 3.3 and $\sqrt{\delta}$ (δ is given by Formula 3.8),

$$P_{conversion,q,d+1} = P_{conversion,q} \min(1, \sqrt{\delta}) \quad (3.17)$$

$P_{click,q,d+1}$ is dependent on the relevant bid which is already generated by bid price generator and stored in knowledge base. We first estimate an exponential function [95] for each query to map bid to position. Then we infer $P_{click,q,d+1}$ according to distributions of click probability and continuation probability provided in the game specification of TAC AA 2009.

3.2.3 Ad Display Selector

Finally, we discuss how to choose an ad display type t_q between *Generic* and *Targeted* for a query q . Generic ad leads to query's system default click-through-rate whereas targeted one can either brings the effective click-through-rate over or under the system default one depending on whether query's component part matches user's underlying component preference. Our following heuristic rule works well in the competition:

$$t_q = \begin{cases} Generic & \text{if } q_c \neq CS \text{ and } q_m \neq MS \\ Targeted & \text{if } q_c = CS \text{ or } q_m = MS \end{cases} \quad (3.18)$$

Based on this rule, for (non-MS,CS) queries or (MS,non-CS) queries, targeted ad will cause a lower-than-default click probability from the users whose underlying product preference

disagrees with our product speciality. However, this is not a truly adverse result. First, users with another component preference are less likely to buy our specialised component. It is pointless to display a generic ad which increases the odds of clicks leading to more cost. Secondly, if a user with different manufacturer preference purchases a product made by our specialised manufacturer, as the manufacturers do not match, our profit is only the standard value. Because the underlying query is a *MS* query, we expect more clicks coming from users with the same preference to purchase *MS* products and yield a larger profit of $USP(1+MSB)$. Since distribution capacity is limited, for users with different manufacturer preference, we would rather like them less likely to click on our ads such that the chance of their low-revenue-same-cost conversion becomes smaller.

3.3 Evaluation

In this section, AstonTAC is analysed from two aspects: competition results to identify the successful properties and three controlled experiments to test robustness of our agent.

3.3.1 Game Results and Analysis

In TAC AA 2009, AstonTAC ranked 2nd out of 15 teams in both qualifying games and the final. In the final, 40 games were played on server one and server two simultaneously. We download logs of all forty games run on server one for analysis in which we are particularly interested in the top three agents - TacTex, AstonTAC and Schlemazl - whose average scores in the final are \$79886, \$76281 and \$75408, respectively.

We start with a correlation test to see whether agent's profit potential is affected by C^{cap} . It turns out correlation coefficient between C^{cap} and average profit is over 97.7%. This proves the importance of adapting bidding strategies to C^{cap} . Such strong correlation also implies that it is only appropriate to make comparison of performance between agents based on same setting of C^{cap} or identical number of each different C^{cap} in case of analysing aggregate results.

As can be seen in Figure 3.3, AstonTAC performs the best in terms of revenue generation. We believe there are two reasons. First, we set high bid prices for high-value queries to target top positions. High-value means high expected profit per click and top position brings maximum number of clicks. Second, we suppress low-profit conversions by bidding less on low-value queries and selecting only profitable keywords to bid on. Moreover, for AstonTAC and TacTex, profit forms a clear ascending trend against capacity whereas for

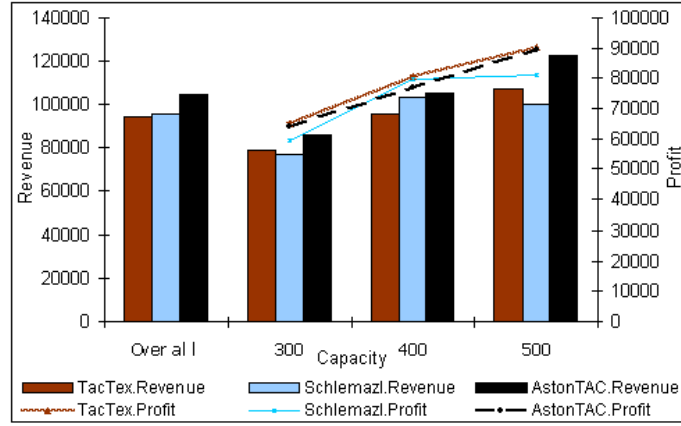


Figure 3.3: Average revenue and profit of top three agents.

The left vertical axis shows the revenue and the right vertical axis shows the profit. The four groups of bars, from left to right, represent the averaged revenue of agents with different capacity: overall (all games), 300, 400 and 500 respectively. The three curves show the correlation between the profit and capacity for each of the three agents.

Schlemazl there no significant profit increase from $C^{cap} = 400$ to $C^{cap} = 500$. The particularly low profit and revenue at $C^{cap} = 500$ indicates that Schlemazl did not sufficiently exploit its conversion space in high capacity. With 25% incremental capacity brought by the change of C^{cap} from 400 to 500, Schlemazl's average number of conversion only increases by 0.96%. By contrast, TacTex's increase rate is 12.3% and AstonTAC's is 14.23%.

CTR is one of the most important criteria to judge whether an on-line advertising campaign is successful. Furthermore, in rank-by-revenue mechanism, high historical CTR can raise rank and reduce payment for the same bid. In TAC AA, high CTR is particularly

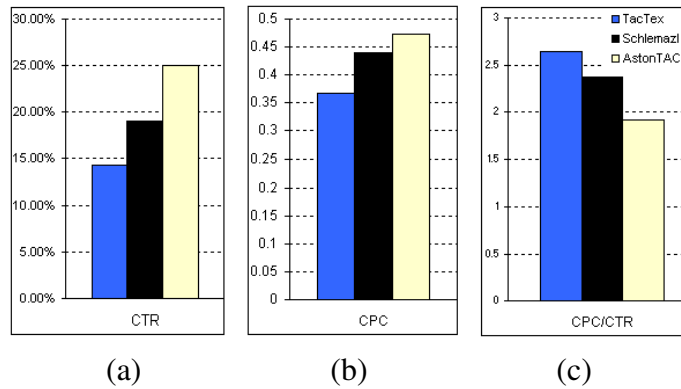


Figure 3.4: (a) Average CTR (b) Average CPC (c) CPC/CTR of AstonTAC, Schlemazl and TacTex.

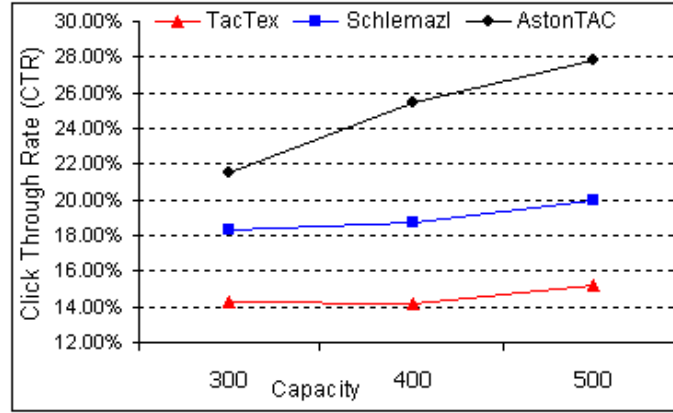


Figure 3.5: Click through rate against capacity.

preferred because parameter χ is closer to 1 than 0 meaning that the ranking mechanism is more by-revenue than by-bid. Based on this, we set high bid prices for high-value queries to target the 1st slot and highest possible CTR. Conversely, our high CTR results in comparatively low CPC. Figure 3.4 shows although our average cost per click is larger than the other two, but our advantage in CTR justifies it. The smallest CPC/CTR means our cost increasing speed against CTR is slower than the other two. To sum up, the rise of cost is dominated by the rise of revenue, our strategy benefits more from increased revenue than suffers from increased cost.

Higher C^{cap} offers the advertiser larger capacity to deal with conversions and subsequently more clicks can be accepted before effective $P_{conversion}$ seriously decays. Therefore increasing CTR against C^{cap} should be expected in successful strategies. Figure 3.5 shows all top three agents have the increasing trend. This trend of AstonTAC is most prominent suggesting its strong adaptivity.

TacTex and AstonTAC are two most successful agents in the competition. But their bidding patterns are distinct from each other as shown by Figure 3.6. We call AstonTAC's pattern surf pattern and TacTex's one jigsaw pattern. Surf pattern makes relatively stable revenue with stable cost every day. In particular, a successful surf pattern agent like AstonTAC can manage to make constantly high profit, meanwhile speculate chances to make even more. Jigsaw pattern makes very high or very low even zero profit intermittently. Other agents' behaviour can also be generally classified into these two categories. Munsey, epflagent typically show a surf pattern and the others show a jigsaw pattern. The reason behind the different patterns is the different ways of treating capacity constraint. Surf pattern aims to maintain a relatively stable amount of conversion everyday to make sure the

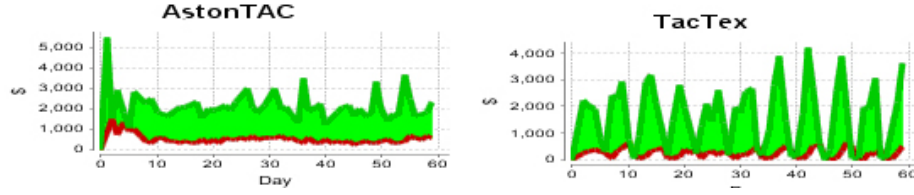


Figure 3.6: Distinct bidding patterns shown by AstonTAC (wave pattern on the left) and TacTex (jigsaw pattern on the right).

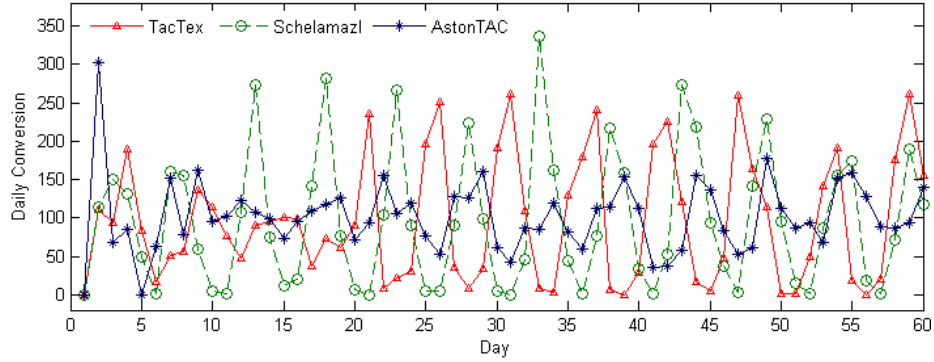


Figure 3.7: Different variances of top three agents in terms of daily conversion.

lowest probability of critical total conversion limit not to be violated and loss not to be suffered. Jigsaw pattern does its best to precisely estimate and sufficiently exploit conversion allowance for any given day. Once it detects that its W days conversion capacity has been completely run out within $W - 1$ days, it stops bidding for one day to let its capacity recover. Even jigsaw pattern finally wins, we still believe that surf pattern has its special merits especially in real world.

After all, our strategy works not because it makes sky-high profit someday but because it consistently makes large profit while other agents could earn nothing periodically. In addition to the second-best profit score, AstonTAC also features lowest-variance in daily conversion. Figure 3.7 demonstrates this through visualizing daily conversions of AstonTAC, TacTex and Schlemazl in a typical final game in which they fortuitously had same C^{cap} and gained very similar accumulative profit. As Figure 3.7 shows, with statistically identical mean, AstonTAC's standard deviation is only 55.9% and 51.2% of that of TacTex and Shlemazl, respectively. This should largely be attributed to behaviour associated with C_{crit} which stretches our ability of making conversions even when C^{cap} is exceeded.

Table 3.2: Settings of post-tournament ad auction games.

Exp.	Participants	Capacity Setting	Iterations
A	AstonTAC, TacTex, Dummy $\times 6$	Game default	80
B	AstonTAC, TacTex, AstonBB, AstonRules, Dummy $\times 4$	Game default	80
C	AstonTAC, TacTex, Dummy $\times 6$	Identical C^{cap}	15

3.3.2 Controlled Experiments

In order to see whether AstonTAC works well in a broader range of environment such as competing with other agents than participants in TAC AA, we have purposefully designed three controlled experiments. Table (3.2) shows settings of each experiment.

Experiment A and B are set in the same way of TAC AA. What is changed is the participating agents. In Experiment C, all agents are assigned identical $C^{cap} \in \{300, 400, 500\}$ in each game.

Experiment A

AstonTAC is the overall winner in Experiment A. In fact, out of 80 games, AstonTAC won 54 whereas TacTex only won the rest 26. Figure 3.8 shows that not only AstonTAC has won more games, but also its winning margin is larger. In particular, there are six games that AstonTAC wins with a margin of over 20000 whereas TacTex has none of these cases. AstonTAC's average score is $50263(\pm 7944.6)$ whereas TacTex got $45439(\pm 8254.9)$. Besides, the deterioration of return-on-investment and cost-per-click comparing with competition results for TacTex are more than that of AstonTAC. We believe our relatively stable performance is in connection with the unpredictable environment caused by dummy agents who are expected to exercise stochastic bidding. Because our bids are based on query's value, AstonTAC is less affected by environmental unpredictability than other agents. In this experiment, both AstonTAC and TacTex's overall performances are worse than that in the final suggesting that the best social welfare can only be achieved when every agent bids wisely.

Experiment B

Two more agents are introduced in Experiment B: AstonBB and AstonRules. AstonBB is initially developed essentially based on balanced bidding strategy [24]. AstonRules em-

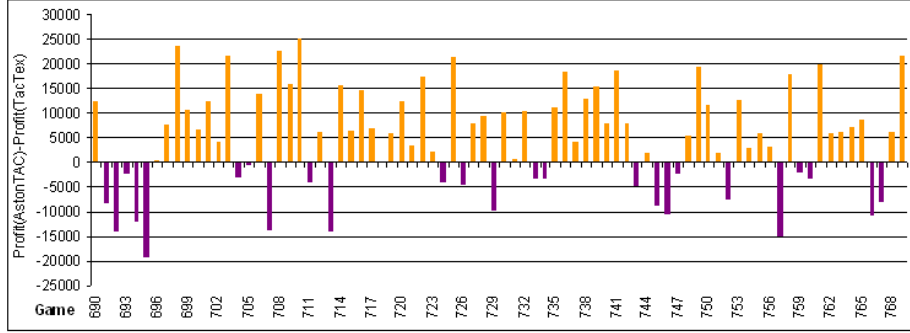


Figure 3.8: Winning margin and frequency comparison for experimental game 690-769. Columns above zero are for AstonTAC, columns below zero are for TacTex (Experiment A).

employs heuristic rules to infer bid according to the position each query receives in each round. AstonTAC is the overall winner again but with a very small margin over TacTex. However, our strategy seems quite superior for high capacity as AstonTAC’s average profit at C_{500} is \$3185 more than TacTex. We believe both *Distribution Capacity Adapter* and *Query Selector* contribute significantly to this result. As for query selector, the algorithm works better at high C^{cap} because prefix decision gets preciser with larger capacity. As the number of intelligent agent increases in the game, TacTex’s performance tends to increase rapidly whereas AstonTAC’s performance keeps stable. It implies that TacTex may have the ability to recognize the bidding pattern of other intelligent agents and act accordingly to undercut the intelligence of their strategy. In contrast, AstonTAC’s strategy is holistically built on the basis of dynamic market-based value per click, which does not need to target any specific position such that it cannot be easily undercut. For this reason, it appears to present stable and reliable performance in whatever environment especially unpredictable ones.

Experiment C

In this experiment, agents performance can be compared directly because capacity bias is eliminated. Out of a total of fifteen games, AstonTAC has won ten (2/3). A table in Appendix D shows the details of each agent’s CTR and profit in each iteration. AstonTAC’s average profit is \$50412 and TacTex’s average profit is \$47269. In most games, we have observed TacTex gradually picks up and forms its jigsaw profit pattern once it found its discipline whereas AstonTAC’s pattern is usually formed after five days. Furthermore, Figure 3.9 shows that TacTex’s click-through-rate forms a declining trend against capacity while AstonTAC’s trend is ascending which is more compatible with intuitions. Our attention

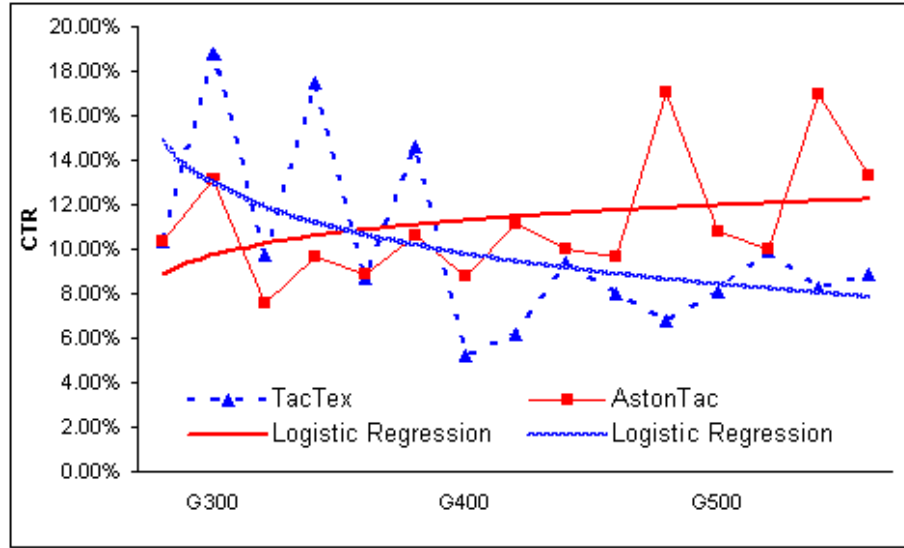


Figure 3.9: Different CTR trend against change of distribution capacity (Experiment C).

to the most crucial game parameter C^{cap} and adaptive bidding strategy designed through different components accordingly should be the reason. For the same reason, this is why AstonTAC's profit increases with capacity is always more significant.

3.4 Conclusion

AstonTAC has shown to be successful and stable across a wide range of settings of TAC AA environments in both the competition and controlled experiments. In particular, we attribute the success of AstonTAC to the strategy used by the bid price generator and the query selector. Market-based Value Per Click reflects the dynamic change of the market and thus leads to the generation of flexible and adaptive bidding prices.

Although strategies employed by AstonTAC are tailored to the specific context of the TAC AA, due to the similar features of the TAC AA and the real sponsored search scenario, many concepts developed for AstonTAC are broadly applicable to an advertiser agent in a real sponsored search scenario. Firstly, the concept of MVPC suits any keyword auction as long as the advertiser aims to make a profit out of clicks after deducting his advertising fees. Secondly, the algorithm of selective bidding suits any advertiser who has a large number of choices on keywords and a constraint whether it is budget or order processing capacity or something else.

In a multi-keyword scenario especially when keywords are also interdependent, it is

essential to view the set of all possible queries as a whole while conceiving a strategy. Selective bidding is one of the ways. MVPC itself is a direct reflection of the market status. Therefore, MVPC-based bid prices are flexible enough to produce a robust performance in almost any circumstance. Given a limited conversion capacity, knowing how much further we can go beyond it gives us the ability of grabbing the largest possible number of conversion to extend our profit space.

3.5 Future Work

As mentioned in Section 3.1, we summarise three research directions in sponsored search: user behaviours modelling, auction mechanism design and advertiser strategy formulation. In the trading agent competition, we have been focusing on the advertiser strategy with the goal of maximising the accumulative profit of the ad campaign. Our MVPC-based adaptive strategy has been proven successful in the specific scenario of TAC AA. Even the real sponsored search *e.g.*, Google Adwords has many different features like some information we can easily obtain from market reports may not be available and the bid price can be changed at any time, the essential ideas in AstonTAC are still useful. When the constraint is changed from distribution capacity to a limited daily budget, our new strategy should consider a maximisation of number of conversions while not exceeding the maximum summed expenditure that we are allowed per day. The combination of heuristic rules and large-scale budget-constraint optimization seems to work well in the real pay-per-click auction like Yahoo Search Marketing. Therefore, we plan to integrate our MVPC-based strategy with Kitts's approach [95].

In addition, we can try different machine learning techniques to generate more accurate estimation of unknown functions *e.g.*, click-conditional-upon-position function and position-conditional-upon-bid function. Revenue-conditional-upon-click-through is one of the most difficult parameter to estimate because of the rarity of sales events. Kitts *et al.* [95] assume that if a customer manages to click through to the site, their probability of generating revenue will be the same after arriving, regardless of the time of their click (*e.g.*, mid night versus midday), or the position from which they click. With this assumption, revenue-per-click can be calculated as a simple average. However, based on the finding of Ghose and Sha [61], the independence assumption does not necessarily hold. Instead, the conversion probability is influenced by the position of the ad on the screen, keyword specific characteristics (retailer, brand and length), the landing page quality score and the time of the click. In our conversion probability model, we may try to include the three

most significant covariants identified in [61]: rank, retailer and brand. Brand information decreases the conversion rate significantly. Retailer information increase the conversion rate significantly. The conversion rate monotonically increases with the rank (position).

Last but not the least, we will design experiments to test the generalisation ability of the three most successful agents in TAC AA 2009: TacTex, AstonTAC and Schelamazl. When other advertisers realise that there is a more profitable bidding strategy, it is rational for them to switch to the better ones. Therefore, it is interesting to see whether the successful strategy in heterogeneous competitive environment can be generalised. We only test the top three strategies because they are most likely to be adopted by other advertisers in a keyword auction campaign. Our question is which of these three strategies generates the highest social welfare. Specifically, we will run homogeneous game in which each of the three top agent is replicated and complete with itself only in every individual game. By comparing the aggregate profit made in each game between games with different homogeneous advertiser population, we hope to reveal which strategy is the best in terms of maximising social outcome and investigate why one strategy generalises better than the other.

Proposal of Multi-item Keyword Auction

Apart from advertiser strategy formulation, even more researchers put their effort in designing a better keyword auction mechanism, which is why the mechanism has evolved from the original generalised-first-price and rank-by-bid to the present generalised-second-price and rank-by-revenue. A new and better mechanism can enhance both the advertisers' utility and the users' satisfaction. Therefore, search engines companies incorporate more and more information into the ranking mechanism in order to induce an automatic behaviour of truth bidding and ad quality improving from the advertisers. Ultimately, better ads improve the user experience and in turn attract more users to click on the ads such that both the advertiser and the search engine can benefit financially from users' clicks.

Billions of people are using the search engine and see the sponsored links everyday, a small improvement of the keyword auction mechanism can make a huge difference. Being aware of the significance of improving auction mechanism, we try to spot problems in the current keyword auction designs that can be improved. Currently advertisers can only bid on and pay for clicks. The publisher chooses the "click" because it is the item that balances the interest and risk of both advertisers and publishers. However, this also means both advertisers and publishers have to compromise. Advertisers have to pay for clicks that do not lead to conversion. Publishers are not paid at all no matter how many times an advertiser's

ad is shown but not clicked. The biggest argument made by the advertiser is: clicks are worthless without conversion. Hence, most advertisers prefer to pay for conversions directly, which is not liked by the publisher because conversion is every difficult to quantify from the publisher side. If search engine can find an effective and impartial way to estimate a user's conversion probability after click-through, offering pay-per-conversion will definitely attract a lot more advertisers. On the other hand, if the goal is brand awareness, it is beneficial to the advertiser to show an advertisement without being clicked. So there is a dispute over the importance of CTR [18, 73]. Some advertisers may not care how the CTR is as long as his ad is shown when users search the keywords of their interest. This is the case especially when the ad link is in the form of logo images or brand names.

After identifying the demand of payment against a variety of items, we propose an the multi-item keyword auction which allows advertisers to freely choose what to pay for. As discussed, there are three different items to pay for in an Internet advertisement: impressions, clicks and conversions. So our new mechanism allows advertisers to bid on any of these three items and positions are allocated according to equivalent pay-per-impression price. In order to convert pay-per-click and pay-per-conversion to a price of pay-per-impression, the publisher will have to estimate both the click-through-rate and conversion probability of an advertiser. Here, we better use an example to explain our idea:

If bidder A bids 5 pence per impression and bidder B bids 200 pence per click and has an estimated CTR of 0.05, B's equivalent pay-per-impression price is $200 \times 0.05 = 10$ such that B will be placed above A. If bidder C bids 20 pounds (2000 pence) per conversion and has an estimated conversion probability of 0.1 and an estimated CTR of 0.06, his bid is equivalent to $2000 \times 0.1 \times 0.06 = 12$ pence per impression, C will be placed above both A and B.

Assuming a default CTR and conversion probability can be obtained based on existing data, every new advertiser's pay-per-click bid or pay-per-conversion bid can be converted into an equivalent pay-per-impression bid for ranking. Existing advertisers' CTR and $\text{Pr}(\text{conversion})$ can be updated hourly mainly based on the statistics of impressions, clicks and conversions. Out of these three metrics, only the conversion data cannot be collected directly by the publisher. So it is designed to be submitted by the advertisers voluntarily. However, the publisher does not need to worry about that advertisers provide false conversion report because it is the advertiser's job to balance between reporting too many conversions and reporting too few conversions. If they report too many, they benefit from increased estimated conversion rate of next round but suffer from paying for every conversion they report. If they report too few, they benefit from saving on the conversions

that they should pay for, but suffer from a much lower estimated conversion rate for next round which significantly increase their cost of bid to stay in the desired positions. Review is made on hourly basis, so eventually, advertise should converge to the rational behaviour – reporting the actual number of conversion. In case that some advertiser do not report conversions at all when they bid on conversions only, some minimum charge depending on the clicks they have received should be applied.

Advertisers in this mechanism have the incentive to improve their ad relevance and quality if they want a position in the sponsored link list because a bad record of CTR or $\text{Pr}(\text{conversion})$ automatically lead their ad to a lower position. There is a way to guarantee a good position regardless of CTR or $\text{Pr}(\text{conversion})$ which is to directly bid for impression. If an advertiser prefer to receive a stable position, he should choose pay-per-impression. If an advertiser prefers a stable cost against fluctuation of CTR, he should choose pay-per-click. If an advertiser does not want to waste money on impressions or clicks, he should go for pay-per-conversion.

In the new regime, advertisers are not restricted to submit a bid on only one of the three items. Instead, they can submit bids on a combination of the three bidding items. For example, a combinatorial bid can be: 1 penny per impression, 10 pence for click and 15 pounds per conversion. Finally, payments are calculated based on the bid profile. For the above bid, if the advertiser receives 100 impressions, 12 clicks and 1 conversion (as reported by the advertiser) in an hour, his payment for that hour will be:

$$1 \times 100 + 10 \times 12 + 1500 \times 1 = 1720(\text{pence})$$

For next hour, this advertiser's CTR and $\text{Pr}(\text{conversion})$ will be reviewed based on the historical data. Consequently, if his bid is not changed, his rank will be calculated based on the updated CTR and $\text{Pr}(\text{conversion})$. The above payment is calculated based on the first-price scheme. We believe it is easy to transfer it to the second-price one if necessary.

To sum up, what is exciting about this multi-item keyword auction is its multi-item feature. Because different advertisers have different demands, it offers them freedom bid and pay for what an advertiser really cares about. Moreover, CTR and $\text{Pr}(\text{conversion})$ are important in deciding an advertiser's rank, which gives the advertiser an incentive to improve the quality their ad, website, products and services.

Chapter 4

Mechanism Design of Double Auction

This chapter focuses on the design of double auction mechanism. Specifically, we will present details of AstonCAT-Plus - a successful e-market specialist designed for and tested in the TAC CAT environment.

After briefly describing the CAT tournament, we will present the details of AstonCAT-Plus. We have carried out experimental analysis in a wide range of CAT environments for different purposes. To evaluate AstonCAT-Plus, two types of experiments are conducted: heterogeneous games and head-to-head games. In heterogeneous games, we compare agent's performance in terms of not only scoring metrics used for CAT competition but also allocative efficiency and convergence coefficient. Through head-to-head games, we show the strength and weakness between two competing market specialists. AstonCAT-Plus performs well (ranked 2nd) in both games, particularly in terms of allocative efficiency and transaction success rates. Through the controlled experiments, our specialist design is shown to be highly effective across a wide range of game profiles.

Apart from the analysis of game results, we have also run trader distribution analysis on heterogeneous games to gain insights about double auction mechanism design. More specifically, we investigate what market mechanism is the most attractive to the traders in a competitive environment. We find that a successful specialist needs to maintain a high market share and a high proportion of intra-marginal traders (see the definition in Sec 4.2.3). Furthermore, we also analyse how the balance between buyers and sellers in a market affects its performance. It turns out that a specialist of market performs well if the *side-balance rate* (see Eq 4.15) is low and with small variance because more trades are produced in such a market. To conclude, a successful market not only needs large trader-population but also balanced trader-profiles.

After investigating the features of an efficient specialist of double auction market, we

proceed to analyse whether the performance of a specialist is affected by the trading strategies employed by the traders. In terms of experimentation, we design seven game profiles. One of the seven represents perfectly balanced trading strategy distribution and the other six either have a dominant strategy or strategy combination. The experimental results show that there does not exist an optimal specialist in the ones that we have tested that can constantly beat all others regardless of the trading strategy distribution; rather, different specialists have their preferred trading strategy (*i.e.*, a specialist performs significantly better if their preferred strategy dominates the market).

In order to help readers understand this chapter and possibly replicate our agent, we provide tables in Appendix C to summarise variables and parameters used in our agent design.

4.1 CAT Market Design Tournament

This section briefly recaps the basic points of CAT. In the CAT tournament, There are two principal entities: trader agents that could be either buyers or sellers, and specialist agents that are double auction market where these traders make deals. All the traders are provided by the tournament organiser, but the specialists are developed by the competition entrants. The platform of the tournament is JCAT [131], a client-server implementation of the Java Auction Simulator API (JASA) [144], which provides additional support for the operation of multiple markets [22]. This platform uses CAT protocol (CATP) detailed in [135] to regulate communication between the CAT server and clients.

Each trader is equipped with a trading strategy and a market selection strategy. The trading strategy determines their bidding behaviour and the decision making process of selecting their offers (or shouts) in a market. The tournament organiser has implemented four most studied bidding strategies of ZI [63], ZIP [35], GD [62], and RE [128]. The market selection strategy specifies traders' behaviour of selecting a market for their trading, which is typically based on their history of making profit in each market. A small number of traders select markets randomly. For those traders who do not randomly select market, they treat the problem as an n-armed bandit problem,¹ and solve it using either an ϵ -greedy exploration policy or a softmax exploration policy [161]. Each trader is endowed with a set of goods to trade and a reservation price (the maximum price willing to purchase for a buyer

¹N-armed bandit problem takes its name from a traditional slot machine where an arm represents a lever. When multiple levers are pulled, each lever provides a reward drawn from a distribution associated with that specific lever.

or the minimum price willing to sell for a seller). Traders can be distributed symmetrically or asymmetrically with respect to their strategies. Based on a pre-set range, traders' private values of reservation price are normally identically, uni-formally distributed and remain constant during each game. However, the specialists cannot know the strategies and private reservation prices of their traders since they are supposed to be private in reality.

The design of a market specialist is to specify the following policies [22]:

- *Shout Accepting Policy*. It determines which shouts are accepted.
- *Clearing Policy*. It determines the way in which bids and asks are matched.
- *Pricing Policy*. It determines the transaction price of matched bids and asks.
- *Charging Policy*. It sets the fees charged to traders and other specialists who wish to use the services provided by the specialist. This kind of fees have the following five types:
 - *Registration fees*: a charge for registering with the specialist.
 - *Information fees*: a charge for receiving market information from the specialist.
 - *Shout Fees*: a charge for successfully placing bids and asks.
 - *Transaction fees*: a flat charge for each successful transaction.
 - *Profit fees*: a share of the profit made by traders, where a trader's profit is calculated as the difference between the shout and transaction price.

A CAT game lasts 500 simulation days in both CAT-2010 and our experiments. Each day consists of 10 trading rounds, each of which lasts for a known constant length of time. At the beginning of each day, specialists announce their fees and traders decide upon which market to register with for that day. Once they make their decisions, the traders cannot switch to another market specialist during that day.

The evaluation of a market specialist is carried out daily against three metrics: (i) *market-share* (i.e., the percentage of the total trader population registered in the market), (ii) *profit-share* (i.e., the ratio of the daily profit obtained by the assessed specialist to the profit of all specialists), and (iii) *transaction success rate* (i.e., the percentage of the shouts accepted that result in transactions). The *daily score* of each specialist is the mean value of these three metrics [22]. The *game score* of a specialist is the sum of their daily scores of 500 simulation days that a game lasts.

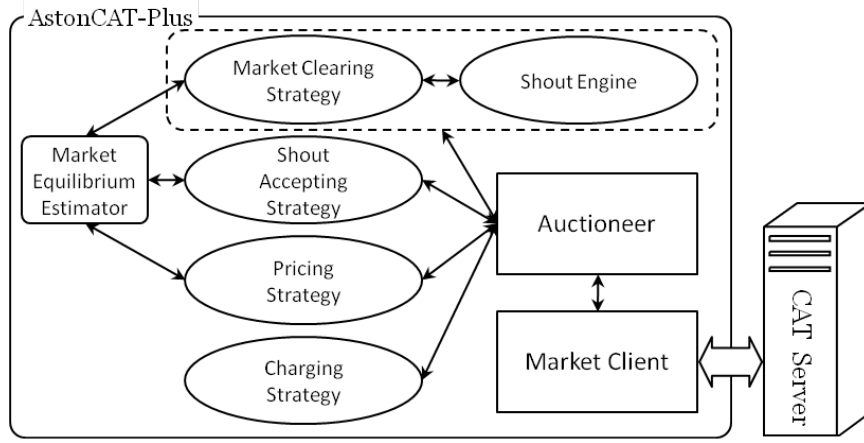


Figure 4.1: AstonCAT-Plus architecture.

4.2 AstonCAT-Plus

This section details the management strategies of AstonCAT-Plus. Figure 4.1 shows its architecture. The four strategies correspond to the clearing policy, shout accepting policy, pricing policy and charging policy, respectively. The *shout engine* registers, sorts and classifies the accepted shouts. It couples tightly with clearing strategy to determine how to match bids with asks. The *auctioneer* acts as a coordinator assembling and passing information requested by other components. The *market client* deals with communication issues with the CAT server. Finally, the *market equilibrium estimator* generates the current estimated equilibrium price, which is referred to by the clearing, accepting and pricing strategies of the market specialist. In this section, we will present their details one by one.

4.2.1 AstonCAT-Plus Equilibrium Estimation

The *equilibrium price* of a market is a price at which the quantity demanded and the quantity supplied are the same [138]. An equilibrium based market mechanism can reduce the fluctuation of transaction prices and achieve a high overall efficiency in their market [133]. So, in order to allocate goods efficiently and price transactions fairly in a market, it is pivotal to find an effective way to estimate its equilibrium price. Our estimation method is based on running sliding windows on two independent streams of market information. One is the history of transaction prices, from which we find short-term equilibrium price p_s . The other is the history of daily maximum transacted asks and the minimum transacted

bids [75], from which we calculate long-term equilibrium price p_l .

When calculating p_s , we use a higher weight for more recent transactions over a short window because p_s is supposed to be reactive to the instant changes of market conditions. Let the last executed transaction be the k th transaction of the game, then p_s is calculated as follows:

$$p_s = \sum_{j=k-W_{short}+1}^k p_t^j \omega_j \quad (4.1)$$

where $W_{short} = 5$ is the size of the sliding window, p_t^j denotes the price of the j th transaction and

$$\omega_j = \frac{0.9^{k-j}}{\sum_{j=k-W_{short}+1}^k 0.9^{k-j}}, \quad (k - W_{short} < j \leq k) \quad (4.2)$$

p_l is actually the rolling average of the middle ground between maximum transacted ask (denoted as \bar{a}) and minimum transacted bid (denoted as \underline{b}). For calculating p_l , we set equal weight on every element over a relatively long window (typically 20) because that way a long-term shifting tendency of market prices can be well reflected. That is, after trading day z is closed, p_l is given by:

$$p_l = \frac{1}{W_{long}} \sum_{i=z-W_{long}+1}^z \frac{\bar{a}_i + \underline{b}_i}{2}, \quad (z \geq W_{long}) \quad (4.3)$$

where W_{long} denotes the sliding window size and when $z < W_{long}$, z itself is used as the window size. Once the value of p_l is obtained, it will be used for the next trading day.

Hence, we can see that p_s contains only a few transactions' information and gets updated dozens of times a day, whereas p_l reflects several days' information and gets updated only once a day. Thus, by combining p_s and p_l , we can obtain a good estimation of equilibrium price, which can balance well the long-term and the short-term market tendencies. That is, the estimated equilibrium price, denoted as \hat{p}^* can be given by:

$$\hat{p}^* = p_s \omega_s + p_l (1 - \omega_s) \quad (4.4)$$

where ω_s is the weight of p_s and the weight of p_l is $(1 - \omega_s)$.

According to the allocative efficiency of 60 heterogeneous games which are run with an ω_s ranging from 0.1 to 0.65 by a step of 0.05, we finally set ω_s to 0.3 because it results the highest mean allocative efficiency. ω_s keeps constant during later experiments.

In fact, the above method of local equilibrium price estimation brings us very satisfactory results. As shown in Table 4.5 (see Section 4.3.1), our transaction prices, which are

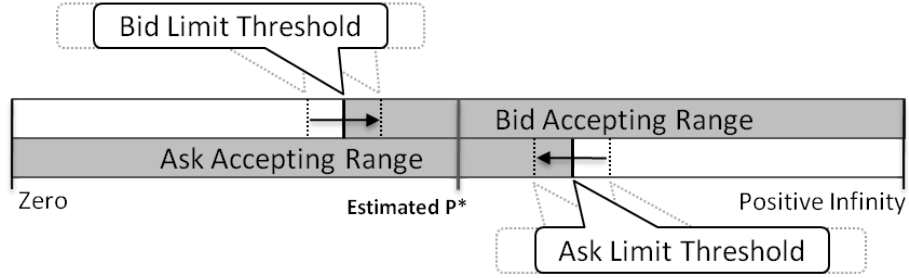


Figure 4.2: Demonstration of AstonCAT-Plus's shout accepting strategy. The shaded areas are accepting ranges. The arrows pointing to the equilibrium price indicate directions towards which accepting thresholds shift during a day.

normally our estimated equilibrium prices according to our pricing strategy, deviate from the theoretical ones by only 6.205 (*i.e.*, 6.28%), which is small comparing with that of other specialists except Mertacor. Moreover, AstonCAT-Plus achieves the highest allocative efficiency meaning that our traders obtain 95.76% of the maximum profit that they can possibly get. Therefore, formula (4.4) and its weight setting are effective for estimating market equilibrium.

4.2.2 Pricing Strategy

Our pricing policy simply sets a transaction price to \hat{p}^* if \hat{p}^* lies inside the bid-ask spread because it implies where demand trades off supply. In the case that \hat{p}^* falls outside the spread, our transaction price is set to the bid or ask price, whichever is closer to \hat{p}^* to avoid a negative transactional profit to one of the two traders. Here the term of *transactional profit* is defined as the profit generated by the difference between bid and ask price. It is different from the concept of *trader profit*, which is defined as transaction price minus seller's private value or buyer's private value minus transaction price. Compared with side-biased pricing policy [177] which basically gives benefit to the under-represented side, equilibrium pricing effectively rewards the intra-marginal (see definition in Section 4.2.3) side in a transaction rather than the side of less quantity. Being short in number does not change an extra-marginal trader to an intra-marginal one in CAT environment.

4.2.3 Shout Accepting Strategy

The shout accepting strategy decides whether a shout can be placed in our market. On one hand, a large shout accepting range can significantly increase the burden on the matching policy [129] and decrease the market efficiency. On the other hand, a small accept-

ing range can rapidly drive potential traders away for the sake of not providing sufficient trading opportunities. Due to the existence of *extra-marginal* sellers (buyers) whose private reservation prices are above (below) the market equilibrium price and *intra-marginal* sellers (buyers) whose private reservation prices are below (above) the equilibrium price, the accepting policy aims to reject unmatchable offers from extra-marginal traders. It is not possible to retrieve the private values directly such that equilibrium price estimation cannot be 100% accurate although we aim to accurately estimate it. For this reason, our accepting thresholds are set around \hat{p}^* (see Figure 4.2) so that the probability of rejecting intra-marginal shouts by mistake can be reduced.

Firstly, it is impossible to estimate the underlying equilibrium price accurately with incomplete information. So our estimated equilibrium price may deviate from the actual one. What's more, with intelligent bidding strategies (*e.g.*, GD, RE and ZIP), intra-marginal traders may attempt extra-marginal shouts to increase their profits. If their shouts cannot get transacted, they will modify the shouts to be more competitive in later rounds. In other words, intra-marginal traders do not always submit intra-marginal shouts in the first place. Therefore, an accepting threshold should have a certain level of tolerance around \hat{p}^* . To this end, we set $\hat{p}^*(1 + \alpha)$ and $\hat{p}^*(1 - \alpha)$ as ask and bid thresholds respectively, where α , called *slack rate*, determines the degree of openness of the accepting policy.

We decrease slack rate α with transactions such that the more the transactions, the tighter the thresholds become. On one hand, a smaller value of α will result in fewer accepted shouts and consequently less transactions than it should be. On the other hand, a larger value of α will result in excess extra-marginal shouts and unfair matches between extra-marginal shouts and intra-marginal ones. Moreover, a too-open policy would reduce transaction success rate due to lots of unmatchable shouts [177]. So, to make a proper trade-off, at the beginning of each day, a large value of α is used to encourage shout submissions, which is important for maintaining market share; and then as transactions are executed, intra-marginal shouts (goods) are consumed, thus the probability of new shouts being submitted by intra-marginal traders and extra-marginal shouts being matched becomes lower and lower, and therefore a decreasing α can effectively block unmatchable shouts from extra-marginal traders and improve transaction success rate. As a result, AstonCAT-Plus achieves the highest transaction success rate in heterogeneous games of the controlled experiments (see Section 4.3.1).

Because a transaction involve at least one intra-marginal shout, the more transactions the less the intra-marginal shouts are in our market. Shifting thresholds towards \hat{p}^* create a easy environment at the beginning to encourage shout submission when intra-marginal

traders have many goods to trade, and make extra-marginal shouts more and more difficult to be accepted when intra-marginal traders have less and less goods to trade.

An ideal initial value, α_0 , of the slack rate is found via experiments between 0.15 and 0.35. Because the number of the buyers and sellers are usually not exactly symmetric, a bias to the side that is inadequately represented would give the fewer side more freedom and result in a more balanced ask and bid profile. The initial value, $\alpha_{0,s}$, of seller's slack rates and the initial value, $\alpha_{0,b}$, of buyer's slack rates need to be updated daily according to the following formulas:

$$\alpha_{0,s} = \alpha_0 \frac{\frac{(n_b+n_s)}{2} + (\beta - 1)n_s}{\beta n_s} \quad (4.5)$$

$$\alpha_{0,b} = \alpha_0 \frac{\frac{(n_b+n_s)}{2} + (\beta - 1)n_b}{\beta n_b} \quad (4.6)$$

where n_s and n_b are the average number of sellers and buyers over last 5 days respectively, $\beta \in [2, 5]$ is used to flatten the result such that the output will not be absurdly far from 1 even if there is a large difference in quantity between buyer and seller. And in the above formulas, actually $\frac{\frac{(n_b+n_s)}{2} + (\beta-1)n_s}{\beta n_s}$ and $\frac{\frac{(n_b+n_s)}{2} + (\beta-1)n_b}{\beta n_b}$ are the *bias ratios* of sellers and buyers, respectively.

During a day, $\alpha_{0,s}$ and $\alpha_{0,b}$ are deducted by a small amount ϵ at every transaction until their pre-defined limits: $l_s = 1.05$ and $l_b = 0.95$. Thus, the ask and bid accepting thresholds τ_s and τ_b can be calculated as follows:

$$\tau_s = \hat{p}^* \max\{1 + (\alpha_{0,s} - n_t\epsilon), l_s\} \quad (4.7)$$

$$\tau_b = \hat{p}^* \min\{1 - (\alpha_{0,b} - n_t\epsilon), l_b\} \quad (4.8)$$

where n_t is the number of transactions happened by the time of the calculation on that day.

4.2.4 Clearing Strategy

After deciding which shouts to be accepted into the market, AstonCAT-Plus will need to decide how to match the accepted shouts and when to convert matches into transactions by using its clearing strategy. AstonCAT-Plus' clearing strategy is a combination of CDA (Continuous Double Auction) and TPT (Transaction Profit per Transaction) based clearing strategy, called *TPT-CDA clearing strategy*. With the CDA scheme, the market is cleared continuously because a transaction takes place as soon as there is a matchable pair of bid and ask. With the TPT clearing scheme, a market is cleared less frequently than with

the CDA one. However, the TPT one still belongs to the category of continuous clearing mechanism compared with *Clearing House*, which clears all matches at one price by the end of a fixed period (e.g., a round).

In our system, the transaction profit per transaction (TPT), denoted ρ , is calculated as follows:

$$\rho = \frac{\sum_{i=1}^{n_{match}} (p_i^b - p_i^s)}{n_{match}} \quad (4.9)$$

where p^b and p^s denote prices of matched bid and ask and n_{match} is the number of matched shout pairs that is equivalent to the number of transactions if they are cleared. Clearly, a trader's transaction profit per transaction is different from its real profit per transaction, which cannot be observed by the market specialist because traders' reservation prices are private. However, the manipulation of transaction profit helps to guarantee deep intra-marginal shouts some minimum real profit so that the chance that deep intra-marginal shouts being exploited by extra-marginal shouts can be effectively reduced, which makes deep intra-marginal traders happier in our market.

Furthermore, from experiments we observe that an adaptive trader has a tendency to offer a price that is close to their private reservation price in order to be competitive in our market. Hence, by promoting traders transaction profit, their real profit is likely to be lifted too as long as they do not pay too much fee out of their transactional profit. Subsequently, if traders make more profit in our market, they will be more likely to return to our market. As for specialist, a high transaction profit per transaction can directly result in a high specialist profit given the same profit fee rate.

With above analysis and speculations, our TPT clearing scheme is designed. It clears matches as long as minimum average transaction profit can be generated from matched shouts. Under TPT scheme, matched shouts pairs within the same set of shouts can be arranged to prevent low-profit or extra-marginal transactions. In practice, we use the TPT clearing scheme only for the first three rounds of a trading day when majority of new shouts are submitted during this period. After the three rounds, trading opportunities are massively reduced as submitted orders have been fulfilled and every trader's trading entitlement is limited. Accordingly, our clearing scheme switches to the CDA one to offer intra-marginal traders the greatest chances to exchange their remaining entitlements within the time left.

A vital issue before clearing is how to match accepted shouts. The two main approaches are: (i) *Equilibrium Matching* (ME) [186], which is commonly used due to its ability of maximising transaction profit; (ii) *Max-volume Matching* (MV) [60] which aims to max-

1. **IF** 'SHOUTPLACED' event occurs **THEN**
2. shout engine sorts matched bid-ask pairs
3. **IF** $round < 3$ **THEN** /* clear market using the TPT clearing scheme */
4. $flag = true$
5. calculate average value, denoted $\tilde{\rho}$, of TPT for matched bid-ask pairs
6. **IF** matched shouts contain extra-marginal ones **THEN** $flag = false$
7. **IF** $flag = true$ **AND** average $\rho > \theta_s$ **THEN** trigger clearing
8. **ELSE IF** $flag = false$ **AND** average $\rho > \theta_l$ **THEN** trigger clearing
9. **ELSE IF** $matching\ volume > \frac{n_{trader}}{10n_{market}}$ **THEN** trigger clearing
10. **ELSE** trigger clearing /* clear market using CDA */

Table 4.1: Pseudo code of AstonCAT-Plus' market clearing strategy.

imise the transaction volume of a market (noticing their procedure does not always produce a maximal matching), but it causes the diversion of intra-marginal traders in a competitive environment because they often suffer profit-loss as being matched with extra-marginal traders. So, we choose ME-based four-heap algorithm [186]² as our shout engine implementation for both TPT and CDA stages. Then, using our clearing strategy, AstonCAT-Plus can offer intra-marginal traders more profit.

Normally all shouts do not arrive at the same time and so their matching problem is dynamic. The traders' entitlements are small and their private reservation prices vary widely, so transaction quality (*i.e.*, transaction profit per transaction) is as important as transaction quantity. Thus, AstonCAT-Plus employs the TPT clearing scheme to improve transaction quality in the first three rounds, and then employs the CDA one to maximise transaction quantity for the remaining rounds. This combination is shown to complement each other and work well to achieve both high transaction quality and quantity in CAT environment. The complete TPT-CDA clearing algorithm is displayed in Figure 4.1. In Section 4.3.1, the average profit of each trader of each specialist are compared to confirm the effectiveness of our TPT-CDA clearing strategy.

Table 4.1 illustrates AstonCAT-Plus' clearing mechanism. There n_{trader} is the total number of traders and n_{market} is the total number of specialists. Two different limits θ_l and θ_s are employed for minimum transaction profit per transaction. Extra-marginal transactions³ do not get executed unless their average value of transaction profit per transaction is

²Four-heap shout engine basically organises shouts into four sorted stacks: (1) Unmatched asks in ascending order, (2) Unmatched bids in descending order, (3) Matched asks in descending order, (4) Matched bids in ascending order.

³Extra-marginal transactions are transactions that involve extra-marginal shouts and the opposite is de-

larger than θ_l . The value of θ_s is considerably smaller than that of θ_l such that intra-marginal transactions have the priority to be cleared unless their profit is too small. However, as matching volume increases, we should not hold the matches for too long because the quantity of transaction is also important in order to maximise traders' total profit in our market. Therefore, the statement on line 9 in Table 4.1 sets a point where the matches are cleared regardless of their transaction profile per transaction. The value of θ_s and θ_l is related to the value distribution of traders' private reservation prices. The smaller the value of the private reservation, the smaller transaction profit per transaction, and therefore smaller θ_s and θ_l ; and vice versa. Assume that seller (buyer) will not attempt an ask (bid) under (over) his private reservation price, the highest attempted bid (\bar{b}_t) and the lowest attempted ask (\underline{a}_t) over a number of days will give an indication of the value distribution of traders' private reservation prices, which confine the maximum value of transaction profit. Accordingly, θ_s and θ_l are set to 2% and 16% of $(\bar{b}_t - \underline{a}_t)$, respectively.

An example of our TPT clearing mechanism is as follows:

- Assuming $\hat{p}^* = 100$, $\theta_s = 2$ and $\theta_l = 16$.
- At time t , suppose matched shouts are: *(bid:120; ask:115)*, unmatched shouts are: *(bid:85, bid:110; ask:125)*. Because matched shouts contain an extra-marginal ask (*ask:115*), θ_l applies based on our algorithm. Transaction profit per transaction $\rho = 120 - 115 = 5 < \theta_l$ and so current match(es) are not cleared.
- At time $t + 1$, suppose a new shout (*ask:105*) is placed, shout engine sorts matched shouts. Matched shouts are: *(bid:120; ask:105)*, unmatched shouts are: *(bid:85, bid:110; ask:115, ask:125)*. Because matched shouts still contain an extra-marginal ask, θ_l applies again. Transaction profile per transaction $\rho = 120 - 105 = 15 < \theta_l$ and so current match(es) are still not cleared.
- At time $t + 2$, suppose another new shout (*ask:90*) is placed, matched and unmatched shouts are re-sorted. Matched shouts are: *(bid:110, bid:120; ask:90, ask:105)*, unmatched shouts are: *(bid:85; ask:115, ask:125)*. For current matched shouts, Transaction profit per transaction $\rho = (30 + 5)/2 = 17.5 > \theta_l$. That is, the clearing condition is met and so current match(es) are now cleared.

defined as intra-marginal transactions. An ask (bid) over (under) estimated equilibrium is identified as extra-marginal shout by AstonCAT-Plus.

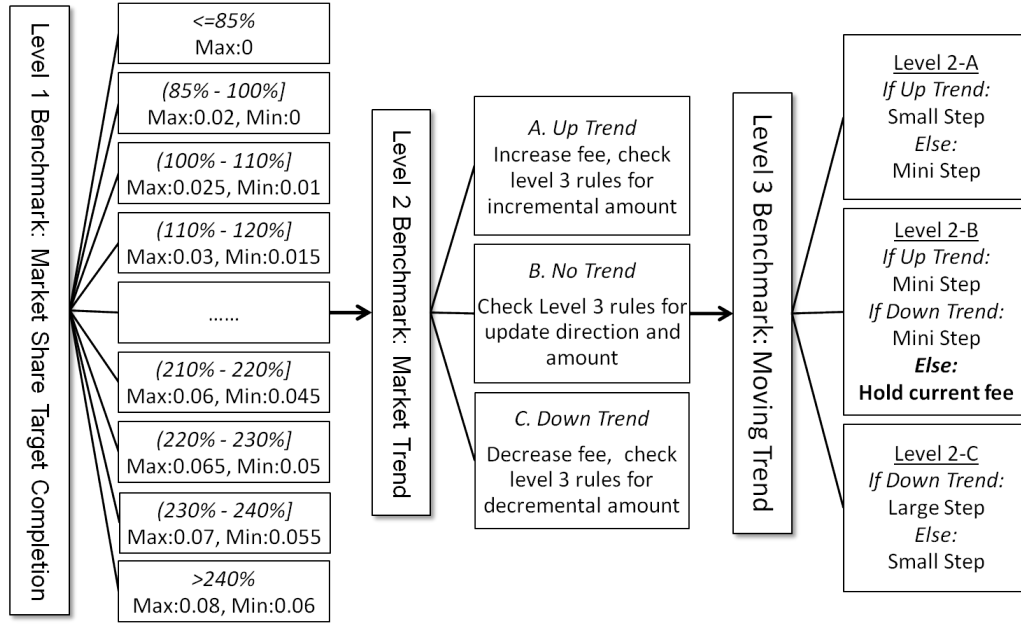


Figure 4.3: Hierarchical market-adaptive stabilized charging strategy.

In the above example, by using our TPT clearing mechanism, we achieve the average value of 17.5 for transaction profit per transaction in our market and the unmatched shouts left are all extra-marginal ones, which is desired. However, if the same shout sequence is cleared immediately, the cleared matches will change to $(bid:120; ask:115)$ and $(bid:110; ask:105)$. Consequently, the transaction profit per transaction becomes $((120 - 115) + (110 - 105))/2 = 5$, which is less than $\frac{1}{3}$ of the average value of transactional profit per transaction that is generated by our TPT mechanism. Furthermore, the remaining unmatched shouts $(bid:85; ask:90, ask:125)$ contain an intra-marginal shout $ask:90$, which is undesired.

4.2.5 Charging Strategy

The charging policy selects the type and the amount of the fees that the registered traders should pay to obtain market services [158]. Out of five fee types explained in Section 4.1, AstonCAT-Plus charges the type of profit fee only. Our experiments show that a free-entry market is more attractive to the traders and thus can lead to a higher market share.

As shown in Figure 4.3, our charging strategy consists of three hierarchical levels of

rules.⁴ The rules at the upper level dominate the rules at the lower ones such that fee updates fired by lower level rules are constrained within a rational range and direction defined by upper levels. Level-1 rules are set based on our current *market share target completion* that is a ratio between AstonCAT-Plus' current trader n_{cur} which is the average trader quantity of last 15 days and trader target, *i.e.*,

$$n_{tar} = \frac{n_{trader}}{n_{market}} \quad (4.10)$$

where n_{trader} is the total number of traders and n_{market} is the total number of e-market specialists in a game. Level-1 functions to confine fees to a rational range instead of updating fees directly. Level-2 determines the direction of fee modification according to *market trend*, which is identified using market trend ratio, *i.e.*,

$$r_t = \frac{n_{cur}}{\bar{n}_{traders}} \quad (4.11)$$

where $\bar{n}_{traders}$ is the all time (From Day 0) mean of AstonCAT-Plus' daily number of traders. When $r_t > 1.16$, it indicates an *up* market trend. When $r_t < 0.92$, it indicates a *down* market trend. And when $0.92 \leq r_t \leq 1.16$, it indicates *no* market trend and thus the decision will be made by using Level-3 rules, which also determines step size of fee updates. The parameter for setting Level-3 rules is called *moving trend* identified by *moving trend ratio* r_v , which is weighted n_{cur} with higher weight for more recent days over n_{cur} . When $r_v > 1.06$, it indicates an *up* moving trend, $r_v < 0.97$ indicates a *down* moving trend, and when $0.97 \leq r_v \leq 1.06$, it indicates *no* moving trend. Combining with decisions made by level-2 rules, appropriate updating step sizes are selected.

Our hierarchical rules have three important properties that correspond to three basic requirements on fees: rationality, adaptivity and stability. Firstly, level-1 rules guarantee that fees are never irrational with respect to a reliable market share evaluation no matter what updates are triggered in other levels. Moreover, level-1 rules make sure our fee range corresponds to the right level according to our market share. Secondly, our fee can be further adapted according to the market trends by using level-2 and level-3 rules. Thirdly, a market trend is identified only if the change is significant. Furthermore, deciding an *up* trend is more cautious to identify an *up* trend than to identify a *down* trend, step sizes chosen to raise fees are smaller than those chosen to reduce fees. All these measures together effectively guarantee the stability of our fees against high volatility of market dynamics. By using the hierarchical rules, the balance between large market share and high profit is

⁴Actually our system can be viewed as a kind of knowledge-based agent system [108].

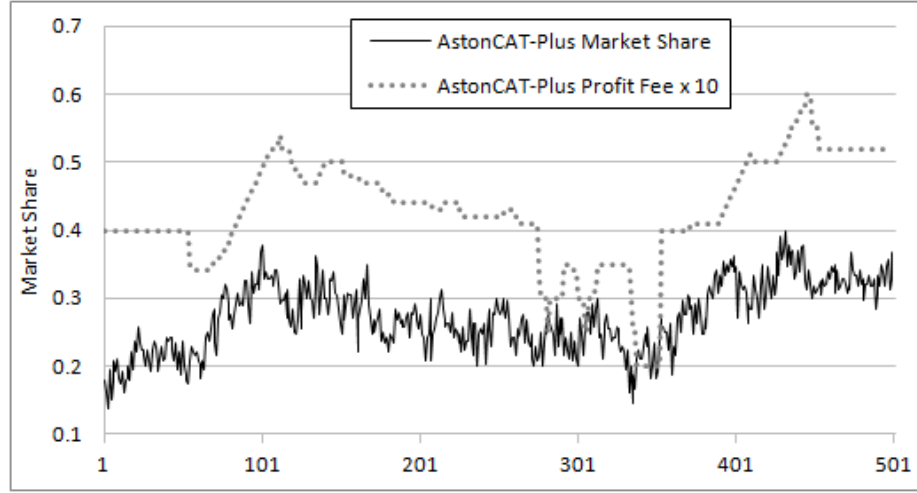


Figure 4.4: Profit fees of AstonCAT-Plus in comparison with its market shares (Data from heterogeneous game iteration 8).

achieved. Figure 4.4 shows the level of profit fee we charge in iteration 4 of heterogeneous games. By comparing with AstonCAT-Plus's market share, we can see both the adaptivity and relative stability of our profit fee.

4.3 Evaluation

This section analyses the performance of AstonCAT-Plus through a variety of controlled experiments and reveals some sights into the design of an effective CAT specialist. Seven specialist agents (see Table 4.2) are included in our experiments.⁵ Three types of experiments are conducted: heterogeneous games, head-to-head games and distribution games. In heterogeneous and distribution games, AstonCAT-Plus competes with five opponent specialists developed by other institutes. In head-to-head games, AstonCAT-Plus competes with only one specialist given in Table 4.2. The first two types of games are similar to the experiments conducted in [132], through which we attempt to test AstonCAT-Plus' performance against its opponents and relate market dynamics to our adaptive auction strategies. The third game explores every specialist's properties with respect to different trader's strategy distribution.

⁵Since AstonCAT-Plus is designed for CAT-2010, we mainly include CAT-2010 agents and the latest version of known successful agents by 2010. The binary code of all agents including AstonCAT and AstonCAT-Plus are available for download at <http://www.sics.se/tac/showagents.php>.

Specialist Name	Description
Mertacor	Winner of CAT-2010
Jackaroo	Runner-up of CAT-2010
TWBB	5th in CAT-2010
PersianCAT	Winner of CAT-2008
IAMwildCAT	CAT-2008 Final
AstonCAT	CAT-2010 Final version of Day 3, ranked the 5th
AstonCAT-Plus	CAT-2010 Post-tournament version

Table 4.2: Specialists used in controlled experiments.

PersianCAT and IAMwildCAT are the latest version available for download at TAC agent repository visited by 30.01.2012. They are chosen due to their outstanding performance in previous CAT tournaments.

Our experiment setting is the same as that of CAT-2010: 500 days each game and 10 rounds each day, and an ϵ -greedy market selection strategy ($\epsilon = 0.1, \alpha = 1$) for each trader. However, all the 500 days are set as scoring days in order to see the performance of each specialist in every stage of a game. Moreover, ten iterations, instead of three in CAT-2010, are run for each game to obtain more statistically significant results. When showing results of an individual game iteration, the representative iteration is selected randomly.

4.3.1 Heterogeneous Games

The heterogeneous games in our experiments are similar to those in the CAT tournament. We set the total number of trading population to 240. Given six specialists in a game, the number of traders per market is set to 40, which is approximate to that of the CAT competition 2010. In order to eliminate any possible bias, the traders are uniformly distributed on the four provided trading strategies of ZI, ZIP, RE, and GD (the default setting of the CAT server). In addition to score metrics, we include allocative efficiency and convergence coefficient to measure efficiency and stability of each market mechanism. We have also introduced some new evaluation metrics like *average trader profit* to reveal exclusive features particularly related to agents' clearing strategies.

Overall Performance

In heterogeneous games, AstonCAT-Plus is ranked the 2nd according to the overall scores (see Figure 4.5), and its average overall score of 10 iterations is only 1.71% lower than Mertacor (winner of CAT-2010) but 29.96% higher than the next best agent Jackaroo

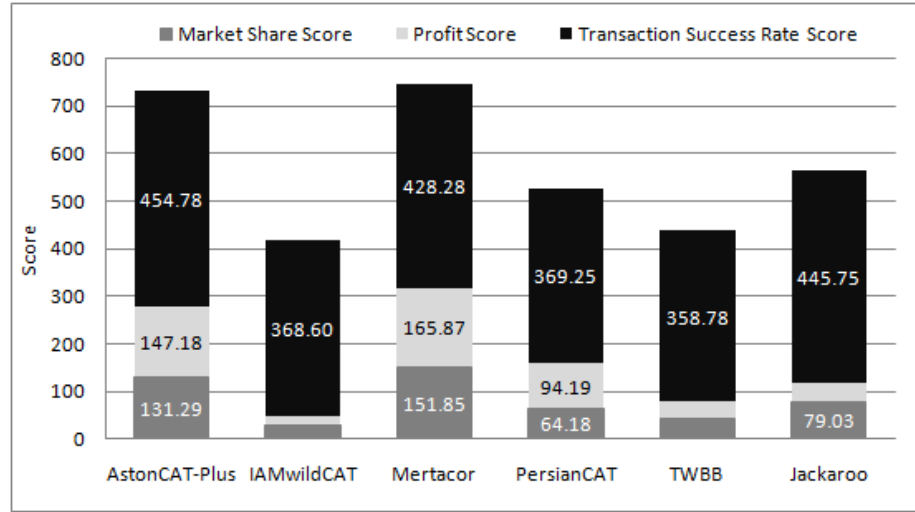


Figure 4.5: Score comparison for heterogeneous games (AstonCAT is not in heterogeneous games).

(runner-up of CAT-2010). Statistically, AstonCAT-Plus' leading over other entrants is significant due to a very small p value ($p\text{-value} \ll 0.0001$) in one tail paired t-test against each of them. Specifically, from Figure 4.5 we can also see that the market share of a specialist is a vital factor for its overall performance as its rank of overall score is monotonically associated with its market share. Hence, undoubtedly the primary target for each specialist is to maximise its market share. On this aspect, Mertacor significantly outperforms every other specialist, while AstonCAT-Plus achieves the second best. And they together dominate the global market because their total market share (56.63%) is considerably more than the total of the rest specialists. A large market share leads to a large number of transactions. By charging a small profit fee (around 4.15% averagely), Mertacor and AstonCAT-Plus made the highest profits too. By comparison, PersianCAT charged 20% profit fee but only obtained 64% and 56.8% of the profit score of AstonCAT-Plus and Mertacor, respectively. So, a small amount of fee charged on a large number of transactions is more effective than to charge high fees, which could sacrifice the transaction quantity of the market in terms of total profit maximisation.

Transaction Success Rate

AstonCAT-Plus achieves the highest *transaction success rate* among all the specialists based on the aggregated results of 10 heterogeneous games. As Table 4.3 shows, it is the only specialist that gained a more than 90% of average transaction success rate through-

Specialist		AstonCAT-Plus	IAMwildCAT	Mertacor	PersianCAT	TWBB	Jackaroo
Transaction Success Rate	Mean	0.910	0.737	0.857	0.739	0.718	0.892
	Stdev	0.0108	0.038	0.0109	0.059	0.098	0.028
	Max	0.920	0.816	0.876	0.797	0.882	0.921
	Min	0.891	0.693	0.841	0.609	0.597	0.824

Table 4.3: Transaction success rates summary for heterogeneous games.

out heterogeneous games. We can also see that AstonCAT-Plus and Jackaroo outperform Mertacor by 4.1% and 6.2% respectively in terms of transaction success rate. Our success on transaction success rate attributes to the shifting threshold accepting strategy, which can block unmatchable extra-marginal shouts submitted by extra-marginal traders outside our market effectively as trading progresses.

The standard deviation of a specialist is calculated based on its daily transaction success rate. According to standard deviations, AstonCAT-Plus and Mertacor facilitate transactions at a stable success rate throughout every game. This gives their traders a stable expectation on chance of transaction, which we believe is one of reasons why traders choose their markets. Moreover, after further analysis of the games, we find that the both markets usually grow from low to high, and stabilise after a period of time. The possible cause is that we both try to favour intra-marginal traders and intra-marginal transactions rather than high transaction success rate. At the beginning of a game, both intra-marginal and extra-marginal traders are arbitrarily distributed in each market, the transaction success rates cannot be very high in order to prevent too many extra-marginal transactions.

Effects of TPT-CDA Clearing Strategy

This section analyses the effects of our clearing strategy, particularly TPT clearing strategy. Our method is to compare each specialist's average trader profit. This is proper because our purpose of introducing TPT clearing strategy is to promote trader profit. So, if AstonCAT-Plus can achieve a high average trader profit, then the new clearing strategy is effective. The average trader profit, denoted $\tilde{\rho}_t$, of a market can be defined as follows:

$$\tilde{\rho}_t = \frac{\tilde{\rho}_r \times \tilde{n}_t}{2} \quad (4.12)$$

where $\tilde{\rho}_r$ denotes the average real profit per transaction of the traders in the market and \tilde{n}_t denotes the average transaction number per trader in the market. In the above formula, there

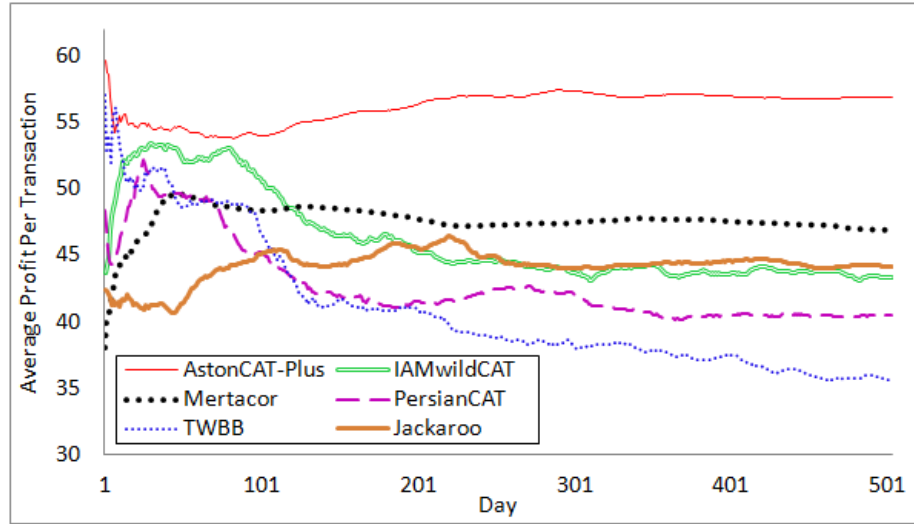


Figure 4.6: Average Trader Profit Per Transaction. Data are from the 8th iteration of the heterogeneous games.

is a denominator 2 because each transaction involves 2 traders: one buyer and one seller. The average real profit per transaction and the average transaction number per trader are not only the components of the average trader profit of a market but also two meaningful criteria for evaluating a market's performance or quality. Therefore, we will compare the markets against these two criteria, respectively.

Figure 4.6 shows AstonCAT-Plus maintains a prominent advantage in terms of average trader profit per transaction. Although its average market share falls behind Mertacor's from day 210, its full-game mean of average trader profit per transaction exceeds that of Mertacor by 9.94 (21.2% of its mean). Regarding daily average trader profit per transaction, only Mertacor and AstonCAT-Plus are stable with the standard deviations 3.72 and 4.30, respectively. This figure of other markets swings violently by a least standard deviation of 9.50. Regarding relative standard deviation of daily average trader profit per transaction to the full-game mean of average trader profit per transaction, AstonCAT-Plus achieves the lowest (7.6%) followed by Mertacor (7.9%). Therefore, AstonCAT-Plus should be particularly favoured by risk-averse traders because it gives them stable and high return expectations.

A high value of average trader profit per transaction of a market alone does not mean that the traders made good profit in the market if their average transaction volume is low. At the beginning of each game, every market's average transaction volume is around 1.5. However, after a short period of time, Mertacor and AstonCAT-Plus establish their lead almost

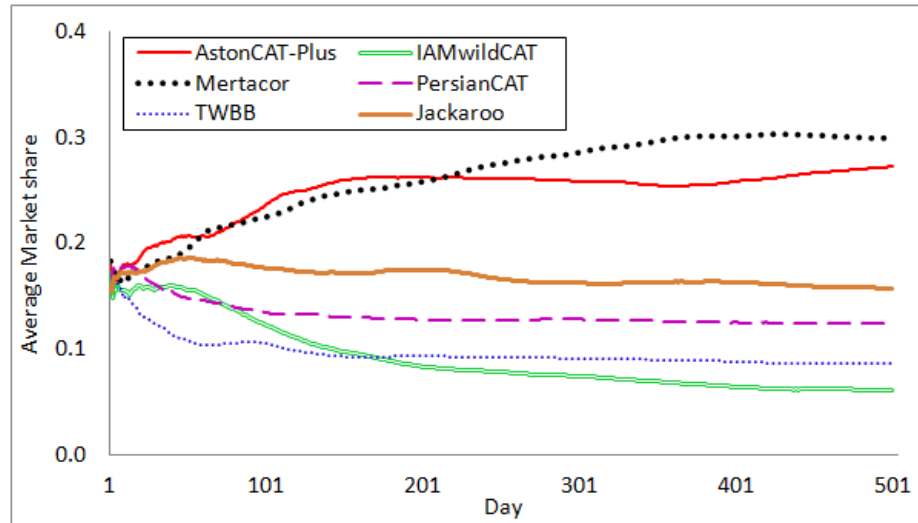


Figure 4.7: Average market share.

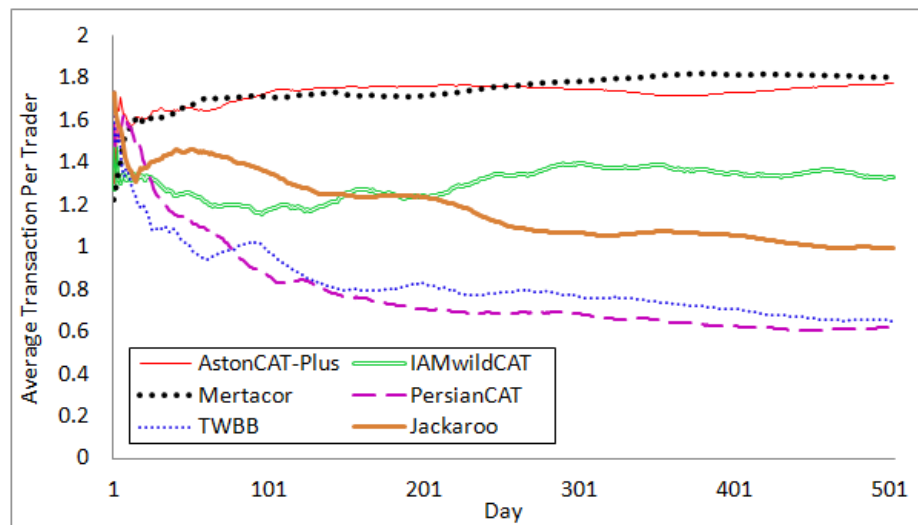


Figure 4.8: Average transaction number (goods traded) per trader.

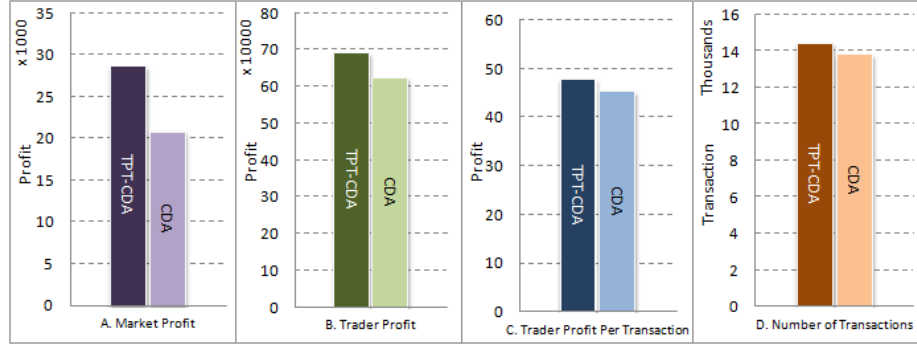


Figure 4.9: Comparison between TPT-CDA and CDA (only) clearing strategy.

	AstonCAT-Plus	IAMwildCAT	Mertacor	PersianCAT	TWBB2010	Jackaroo
Mean	47.593	30.722	44.472	11.856	15.528	22.392
Stdev	3.538	3.368	3.161	2.505	5.383	2.637

Table 4.4: Mean and standard deviation of average profit per trader of across 10 heterogeneous game iterations.

simultaneously. Figure 4.8 shows that at the end of the game, the traders with AstonCAT-Plus have traded averagely 1.77 goods by average out of three total entitlements, which are 33.5%, 187.2%, 172.3%, 78.4% more than IAMwildCAT, PersianCAT, TWBB, Jackaroo respectively, but only 1.6% less than Mertacor. AstonCAT-Plus do not seem to suffer from reduced transaction volume while trying to achieve high transaction profit. In clearing strategy, two things are done to enhance transaction volume. (1) CDA is employed after 3 rounds. (2) In TPT, we encourages intra-marginal transactions using a smaller transaction profit threshold.

Finally, the average profit per trader of AstonCAT-Plus is 50.38 in iteration 8, which significantly exceeds those of Mertacor, IAMwildCAT, Jackaroo, PersianCAT and TWBB by 19%, 75%, 130%, 303% and 335%, respectively. It is noticeable that IAMwildCAT's good average profit per transaction does not seem compitable with its poorest overall performance. This is probably due to a narrow and rigid shout accepting range that guarantees good transaction experience for accepted shouts from highly intra-marginal traders but drives away a large number of average intra-marginal traders at the same time. Table 4.4 shows the overall performance across all 10 heterogeneous game iterations in terms of average profit per trader. AstonCAT-Plus achieves the best (47.593) by this criteria despite a slightly big variance.

Besides the positive results shown by heterogeneous games, we have done additional

Specialist		AstonCAT-Plus	IAMwildCAT	Mertacor	PersianCAT	TWBB	Jackaroo
Allocative Efficiency %	Mean	95.762	82.744	95.474	71.625	68.301	87.152
	Stdev	0.541	2.316	0.239	4.120	8.598	2.305
	Max	96.539	87.175	95.800	75.944	83.540	90.681
	Min	94.480	79.396	95.116	63.580	57.780	82.553
Convergence Coefficient	Mean	6.205	12.502	5.163	8.013	10.308	8.065
	Stdev	0.816	1.211	0.722	1.221	0.749	0.980
	Max	7.217	14.114	6.061	9.506	11.592	9.266
	Min	4.896	10.702	3.943	5.844	8.889	6.632

Table 4.5: Summary of allocative efficiency and convergence coefficient in heterogeneous games.

experiments specifically designed to test the effectiveness of our TPT-CDA clearing strategy. In these experiments, we just replace TPT-CDA clearing strategy with only CDA clearing strategy and run 10 iterations of each setting. Figure 4.9 shows the average result of 10 game iterations in terms of market profit, the total trader profit, the average trader profit by number of transactions and the number of transactions. First of all, we can see a significant improvement of market profit, which is easy to understand because market profit directly benefit from lifted transaction profit. Secondly, the trader profit with respect to the CDA clearing scheme falls by 9.6% in comparison with that of the TPT-CDA one and the average trader profit per transaction falls by 4.3%. Hence, the average number of transactions of TPT-CDA market also beats CDA market by 14451 vs 13830. Furthermore, AstonCAT-Plus' allocative efficiency dropped by 3.23% when its clearing scheme is switched from the TPT-CDA to pure CDA. So, we are confident that the TPT-CDA clearing scheme is an effective and efficient at least for AstonCAT-Plus in the context of CAT.

Efficiency and Convergence

Allocative efficiency and convergence coefficient are two essential metrics to identify whether or not a market is efficient and stable. According to [22], the *allocative efficiency* denoted as φ is defined as the ratio of the traders' actual profit to the theoretical maximum profit (obtained had all traders traded at the theoretical equilibrium according to microeconomic theory [143]):

$$\varphi = \frac{(\sum_{i=1}^n v_{b_i} - p_i) + (\sum_{j=1}^m p_j - v_{s_j})}{(\sum_{i=1}^n v_{b_i} - p_0) + (\sum_{j=1}^m p_0 - v_{s_j})} \quad (4.13)$$



Figure 4.10: Allocative Efficiency of each agent in a typical heterogeneous game.

where p_0 is the theoretical equilibrium price; v_{b_i} and v_{s_j} are buyer i 's private value and seller j 's private value, respectively; p_i and p_j are buyer i 's actual transaction price and seller j 's actual transaction price, respectively; n and m are buyer's number and seller's number, respectively. The *convergence coefficient* denoted as ψ is defined as the standard deviation of transaction prices around daily theoretical equilibrium:

$$\psi = \frac{100}{p_0} \sqrt{\frac{1}{n_t} \sum_{i=1}^{n_t} (p_{a,i} - p_0)^2} \quad (4.14)$$

where n_t denotes the number of transaction and $p_{a,i}$ denotes the actual trade price of transaction i . According to the heterogeneous game results shown in Table 4.5 and Figure 4.10, Mertacor and AstonCAT-Plus' efficiencies are significantly higher with significantly smaller standard deviation than other specialists, which means they are much more efficient and stable markets.

Figure 4.11 shows convergence coefficient of each specialist in a typical heterogeneous game. For days convergence coefficient cannot be calculated in some markets, it is denoted by a negative number of -10. We can see only Mertacor, AstonCAT-Plus and Jackaroo do not have a day that equilibrium cannot be found. Mertacor got the most steady convergence coefficient, which is almost always under 10. AstonCAT-Plus keeps convergence coefficient under 20 constantly. Jackaroo occasionally got it over 20. With highest convergence coefficient nearly 80, IAMwildCAT is worst agent measured by this metric, which is probably the main reason for its poorest overall performance. Low convergence coefficient often comes with high allocative efficiency because market is often cleared at somewhere near the equilibrium price. Since our pricing policy prices transactions at the estimated market equilibrium price \hat{p}^* , the second lowest aggregated convergence coefficient according to Table 4.5 demonstrates the effectiveness of our method for equilibrium estimation. Mertacor's lowest convergence coefficient (5.163 ± 0.722) indicates that its equilibrium estimates are more accurate, which is probably why AstonCAT-Plus cannot beat it.

4.3.2 Head-to-Head Games

In heterogeneous games, the performance of any market specialist depends not only on its own policy but also on the policies of other competitors. Head-to-head games make direct comparisons between AstonCAT-Plus and its opponents to identify its strength and weakness in such environment. In head-to-head games, without loss of generality the trader population is set to 120 and the four provided trading strategies (see Section 4.1) are evenly

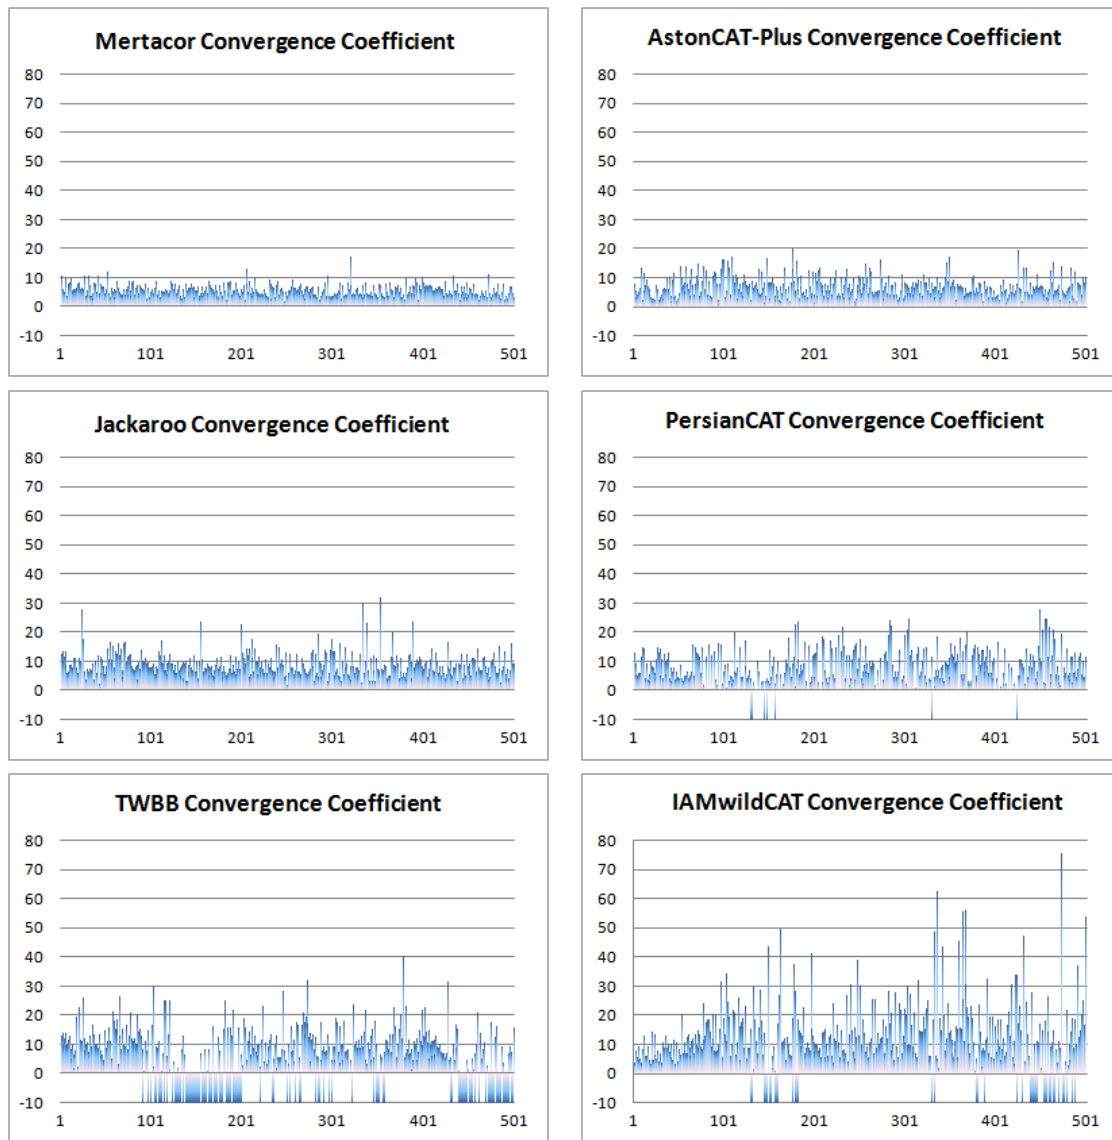


Figure 4.11: Convergence coefficient of each agent in a typical heterogeneous game.

Opponent	Overall Score	Market Share	TSR	Efficiency %
IAMwildCAT	0.842 vs 0.402	0.774 vs 0.226	0.862 vs 0.869	93.02 vs 93.41
Mertacor	0.558 vs 0.659	0.462 vs 0.538	0.839 vs 0.824	92.08 vs 94.57
PersianCAT	0.757 vs 0.475	0.693 vs 0.307	0.908 vs 0.789	94.08 vs 70.29
TWBB	0.789 vs 0.470	0.750 vs 0.250	0.884 vs 0.893	94.05 vs 82.77
Jackaroo	0.681 vs 0.574	0.574 vs 0.426	0.836 vs 0.930	94.32 vs 92.85
AstonCAT	0.853 vs 0.356	0.663 vs 0.337	0.898 vs 0.730	94.11 vs 81.39

Table 4.6: The results of head-to-head games.

Each repeated 10 times. First values in each column refer to mean of AstonCAT-Plus and second ones refer to means of the corresponding opponents.

distributed over the trader population.

Overall Performance

The overall results are consistent with that of heterogeneous games (see Table 4.6). Mertacor is still the only market specialist that outscores AstonCAT-Plus. AstonCAT-Plus' market shares are very stable in head-to-head games. Our smallest market share 46.2% comes from the games played with Mertacor. Although our transaction success rates are outscored by IAMwildCAT, TWBB and Jackaroo, we still obtain significantly more market shares against each of them. With further investigation, we realised Mertacor's transaction success rate of first 180 days are much lower than those of the remaining 320 days, which indicates that its transaction success rate is sacrificed for the satisfaction of intra-marginal traders at the beginning of the game. This strategy certainly worked according to Mertacor's outstanding overall performance and stably high market share in any circumstances. In addition, its high transaction success rate for the rest of the game compensates for the loss of transaction success rate in the first 180 days, which brings its overall score of transaction success rate to an acceptable level in the end. Thus, it is suboptimal trying to maximise transaction success rate in CAT before an intra-marginal trader dominated market is established. To conclude, AstonCAT-Plus' allocative efficiency is high and stable (varying between 92.08% and 94.32%) despite opponents, which shows the robustness and reliability of our auction mechanism.

Market Convergence Trend

In the head-to-head situation, we are also interested in whether or not the traders converge to one of the two markets. If so, how the migration and convergence take place. Fig-

ure 4.12 shows that the market quickly converges to AstonCAT-Plus in the games against PersianCAT, TWBB and IAMwildCAT, gradually converges toward AstonCAT-Plus in the games against AstonCAT and Jackaroo. In the game against Mertacor, AstonCAT-Plus managed to hold an equilibrium of market share where traders do not converge to either market.

AstonCAT-Plus vs AstonCAT

The original AstonCAT showed a significant improvement on the third day of the CAT-2010 because it employed the framework that is similar to AstonCAT-Plus, which accepts shouts and clears matches based on estimated equilibrium of the market. The improvement is not only reflected by the better rank (5th in Game 3 vs 8th in both Game 1 and 2) and higher score (176.12 in day 3 vs 120.84 in day 1 and 130.71 in day 2), but the second best efficiency (94.17%), which exceeds that of day 1 by 55.8%. Although their frameworks are similar, AstonCAT-Plus is more sophisticated and accomplished on many aspects. For example, the equilibrium estimator of AstonCAT is based on transaction prices only; the accepting thresholds of AstonCAT are fixed throughout a game; and AstonCAT's charging strategy is not systematic such that it could be vulnerable to certain composition of entries.

Rather surprisingly, AstonCAT-Plus' overall score is a massive 240% of that of AstonCAT. So, obviously the current version has successfully sorted out the major problems of the original version. By further investigation, we realise that AstonCAT's charging strategy leads to unstable profit dispersion because whether or not to enforce fees depends too much on its competitors' charging policies.

AstonCAT-Plus is not only better in score but also stabler and less vulnerable than AstonCAT according to their relative performance in head-to-head games. First, AstonCAT-Plus's average overall score is 239% of that of AstonCAT. Specifically, AstonCAT-Plus outperforms AstonCAT by 96.2% on market share, 23.0% on transaction success rate, and 15.6% on allocative efficiency, respectively. The biggest difference between them is in their profit scores. AstonCAT-Plus nearly obtains full profit score (499.2 vs 0.8). The reason is that AstonCAT-Plus can maintain a high market share persistently, which makes AstonCAT's market share too low to allow its minimum charging rules to be fired. In other words, a free-of-charge market is not necessarily competitive in attracting traders against a charged market. This indicates that to improve traders' profitability is more effective than to lower fees in keeping traders.

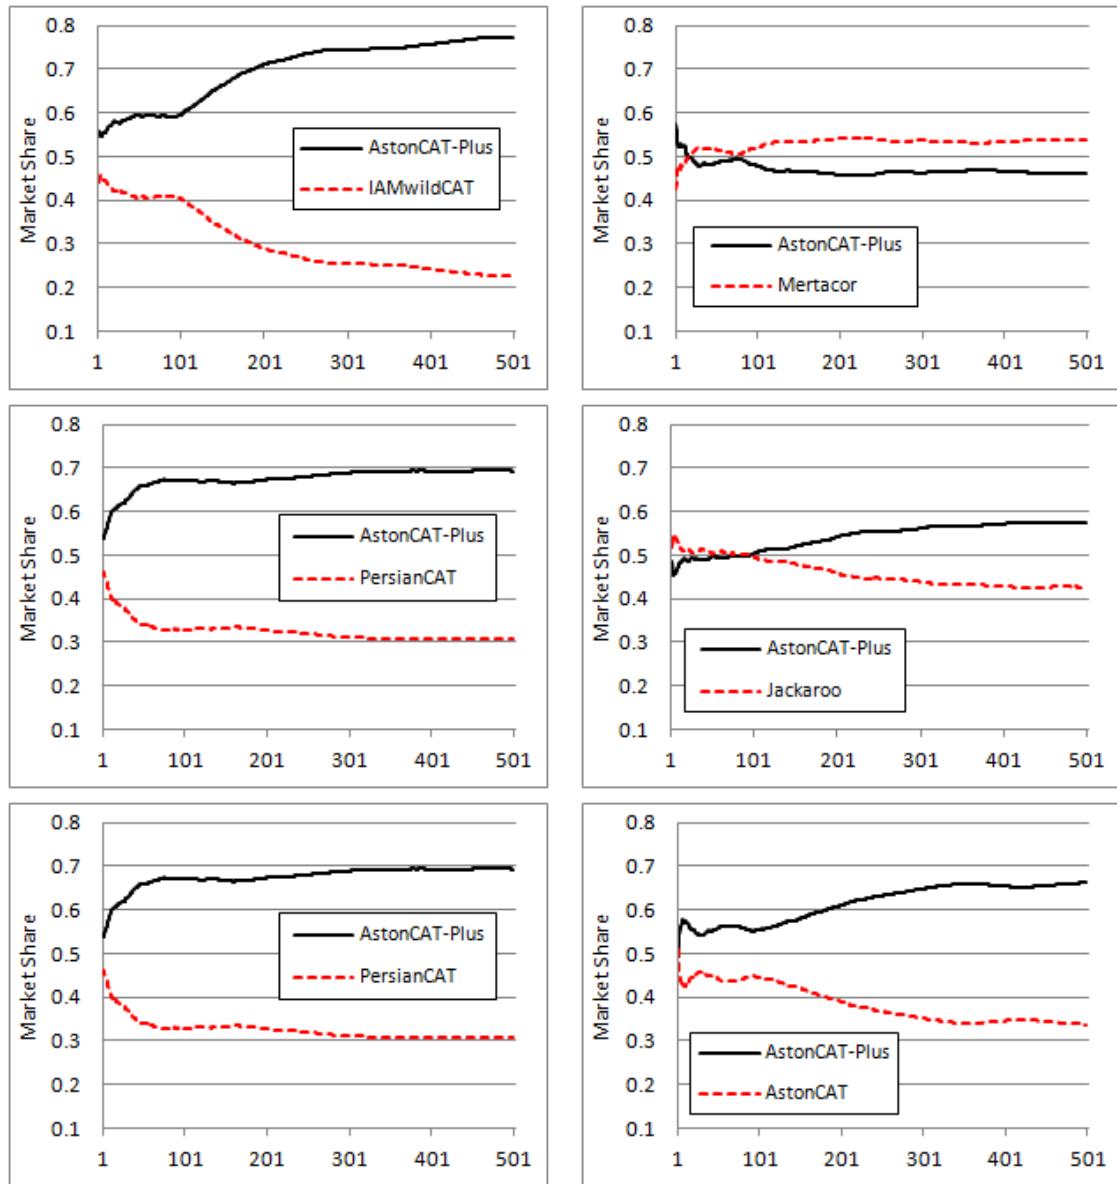


Figure 4.12: Daily market shares in head-to-head games.

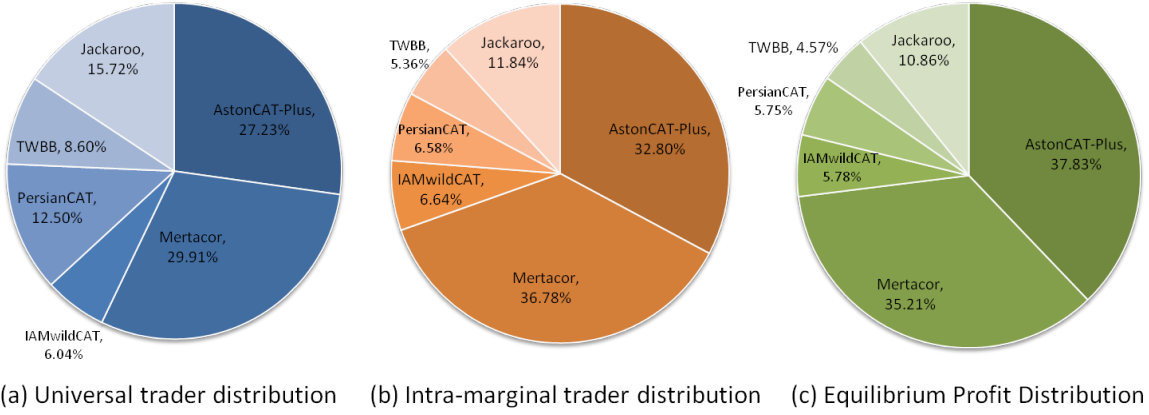


Figure 4.13: Distributions of universal traders, intra-marginal traders and equilibrium trader profits in a typical heterogeneous game.

4.3.3 Trader Distribution Analysis

One of the main objectives of CAT tournament is to investigate which market mechanism is the most attractive to traders in a competitive environment of multiple markets. Hence, based on the data of a representative heterogeneous game, this section analyses the distributions of traders (especially intra-marginal traders) and trader reservation prices to reveal some important properties of AstonCAT-Plus and other CAT specialists.

Trader Distribution

Figure 4.13 demonstrates three distributions: (i) the universal trader distribution, which is equivalent to the market share of each specialist, (ii) the intra-marginal trader distribution, which shows the market share of global intra-marginal traders,⁶ and (iii) the equilibrium profit distribution, which shows the percentage of each market's theoretical trader profit to aggregate theoretical profit of all markets. Successful specialists like the winner of CAT-2010 and our AstonCAT-Plus have the same feature, which is the ability to keep far larger than average market share. Furthermore, in terms of the shares of intra-marginal trader, their dominance is even more prominent.

Each specialist's average market share should be $\frac{1}{6}$ (*i.e.*, 16.7%) when the traders select markets randomly. As shown by Figure 4.13(a), with respect to the universal trader distribution, Mertacor's and AstonCAT-Plus largely over-fulfilled their average market share by 13.3% and 10.5%, respectively. Consequently, none of the other specialists did manage to

⁶Global intra-marginal traders are those whose private values are superior to global equilibrium.

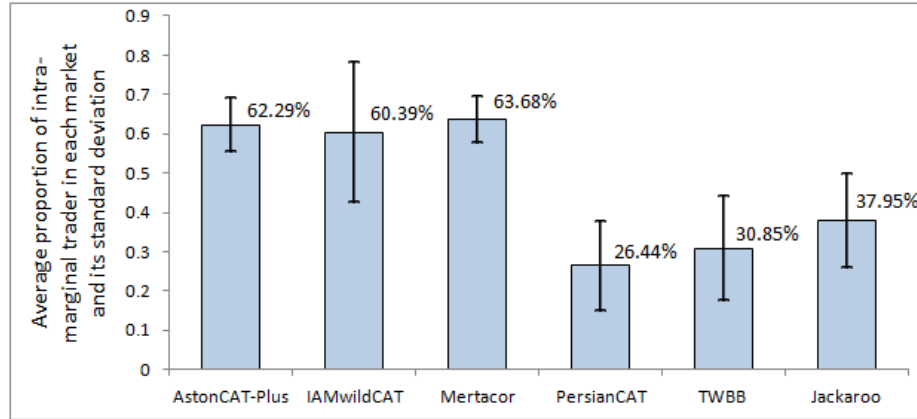


Figure 4.14: Mean and standard deviation of percentage of intra-marginal trader to total number of traders.

achieve this in the experiments.

In fact, the private value of each trader is available in the game data after a game is completed. According to the data, the intra-marginal traders can be identified and the percentage of the intra-marginal traders can be calculated for each specialist of market. As far as the distribution of global intra-marginal traders is concerned, Mertacor and AstonCAT-Plus extended their dominance by a big margin. Their aggregate market share is increased to 69.6% from 57.1%. This means that the trend of convergence of intra-marginal traders into these two markets are even more prominent as shown in Figure 4.13(b).

The intra-marginal trader can be further divided into: (i) deep intra-marginal trader whose private reservation price is greatly superior to the global theoretical equilibrium, and (ii) shallow intra-marginal trader whose private value is slightly superior to the global theoretical equilibrium. For the same number of intra-marginal traders, the deeper the intra-marginal traders the better because deeper ones can produce higher profit to themselves and their trading counterparts. Figure 4.13(c) shows the percentage of profit of each market in terms of trader values. Generally, the shares of equilibrium profit agree with the shares calculated based on the number of intra-marginal traders. However, the biggest difference between Figures 4.13(c) and (b) is that AstonCAT-Plus' share is increased by 5.03%, while every other's is decreased as shown in Figure 4.13(c) compared with their values in Figure 4.13(b). Apparently, AstonCAT-Plus is competitive on keeping deep intra-marginal traders. We attribute this achievement to our clearing strategy that highly rewards deep intra-marginal traders using larger threshold and effectively facilitates transactions between shallow intra-marginal traders using smaller threshold.

Proportion of Intra-marginal Traders

Here we focus on exploring the internal constitution of traders within each specialist. Figure 4.14 shows the mean and standard deviation of the intra-marginal trader proportion in each specialist market. As shown, every market contains both intra-marginal and extra-marginal traders. Mertacor and AstonCAT-Plus maintain a high proportion of intra-marginal traders with the smallest standard deviations (*i.e.*, AstonCAT-Plus: 6.9%; and Mertacor: 6.0%). IAMwildCAT's average intra-marginal trader proportion is also high (60.39%), but its standard deviation is the highest among all specialists. According to trader distribution analysis, IAMwildCAT's average market share is around the lowest one. We believe this is the main reason that causes its unstable trader structure. Hence, another feature of a successful specialist is a high and consistent intra-marginal trader composition. Another interesting finding is that it is difficult to eliminate all the extra-marginal traders in one market because the highest average intra-marginal trader proportion is just over 63%. However, it is not hard to understand this because the more intra-marginal traders in a market, the more attractive the market is to extra-marginal traders.

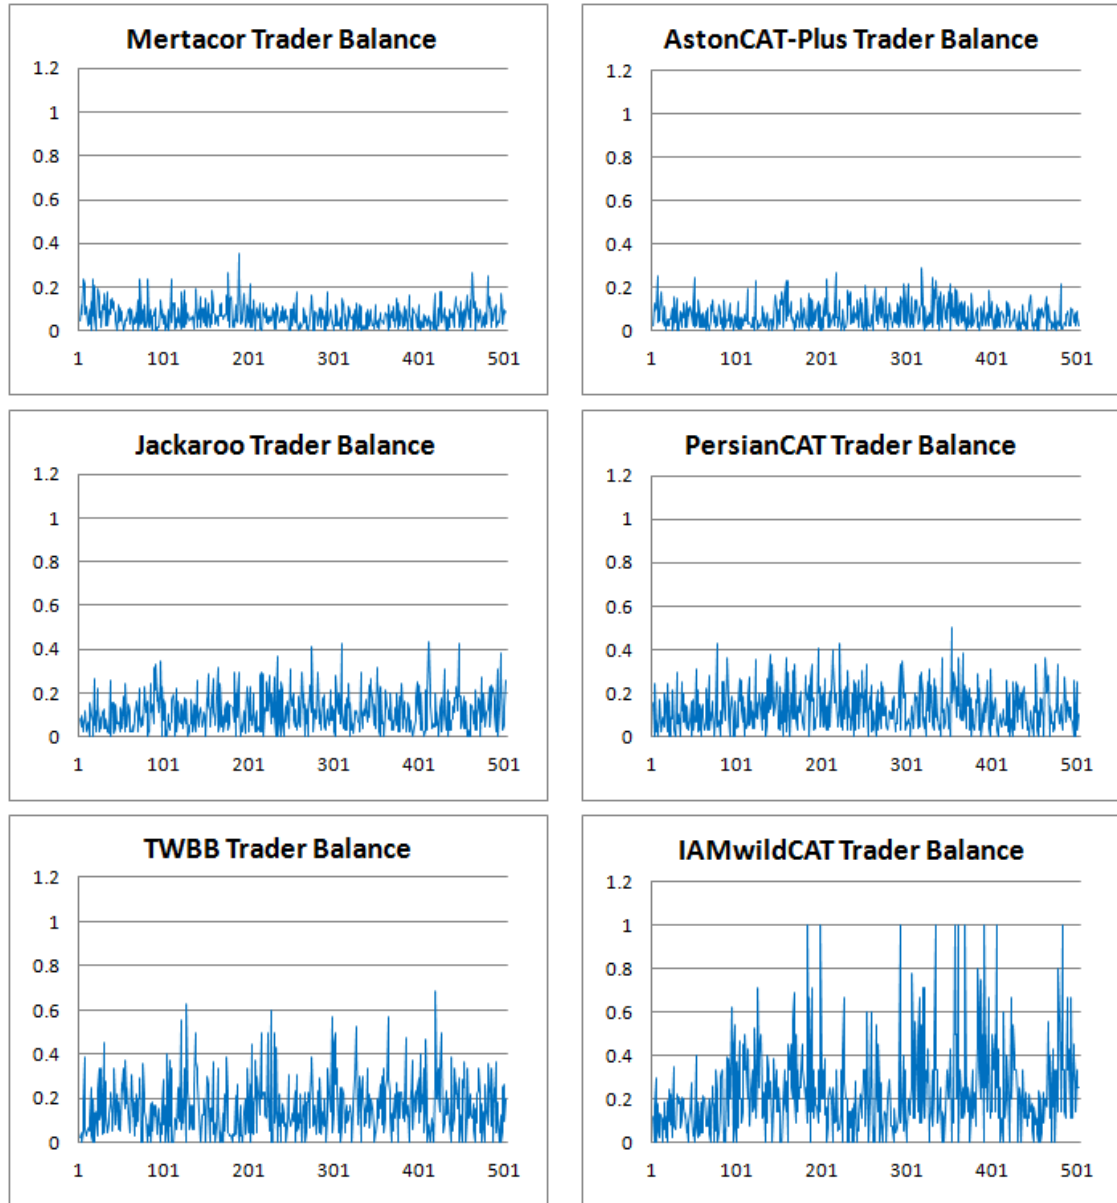
Trader Balance

A well-formed double auction market should maintain a good balance between supply and demand. In CAT tournament since every trader has the same amount of goods to trade, the balance between supply and demand is equivalent to the balance between the seller and the buyer. When the trader balance is lost, the less represented side is compensated with a looser shout accepting threshold in AstonCAT-Plus market in order to create a relatively balanced accepted shout profile and facilitate as many trades as possible (see Section 4.2.3). Thus, our mechanism can result in a desired balance between demand and supply, which is consistent with the symmetric and fixed global demand and supply.

To measure the balance of demand and supply in a specialist market, we introduce side-balance rate λ that is defined as follows:

$$\lambda = \frac{|n_s - n_b|}{n_s + n_b} \quad (4.15)$$

where n_s and n_b are the numbers of sellers and buyers in a market on each day. A market mechanism that performs well (*i.e.*, with a smaller value of λ) is capable of facilitating more trades. If a trader finds that he can trade more goods in market A than in market B , the trader is more likely to prefer A to B . So, we expect successful specialists to

Figure 4.15: Trader side-balance rate λ of each specialist.

Game Type	Trader structure	ZIP	RE	GD	ZI	Total
I	Distribution A	120	120	60	60	360
	Distribution B	90	90	90	90	360
	Distribution C	60	60	120	120	360
II	Distribution ZIP	280	40	40	40	400
	Distribution RE	40	280	40	40	400
	Distribution GD	40	40	280	40	400
	Distribution ZI	40	40	40	280	400

Table 4.7: Trader population settings for strategy distribution games.

keep this rate as close to zero as possible. Top three specialists Mertacor, AstonCAT-Plus and Jackaroo all have consistently small values of λ . In particular, AstonCAT-Plus and Mertacor both can keep λ under 0.2 during almost the whole game in addition to far smaller variances compared with the other specialists. This evidently reflects the fact that a successful specialist has the feature of a stably balanced trader profile in the competitive global environment. Figure 4.15 shows the traders side-balance rate for each specialist.

4.3.4 Trading Strategy Preferences

In order to test the robustness of our results in the setting that is different from that of CAT-2010, we have conducted some special experiments where each type of games is featured by a different trader structure with respect to trading strategy distribution. Two types of experiments were carried out. In type I, distribution A and C simulate the trader structures that are similar to that of [132] where two of the four provided strategies are over-represented while the other two are under-represented. And distribution B produces a uniform distribution of strategies. In type II, there is one dominant trading strategy, which equips the majority of the trader population. With this type of games, we aim to explore exactly which trading strategy is preferred by which specialist if such a preference exists. We expect that a specialist performs well when its favourite trading strategy is dominant. There are seven individual game profiles. Table 4.7 shows the experiment set-up for each game profile which mainly specifies the number of traders using each different trading strategy. With respect to each profile, we have run five heterogeneous games.

Figure 4.16 displays the average daily scores of each specialist in type I. AstonCAT-Plus maintains its rank against the change of trader structure. Its smallest advantage over the third is 7.7% across distributions A , B and C . This demonstrates its robustness. The relative rankings between each market agent keep unchanged across three games of type I

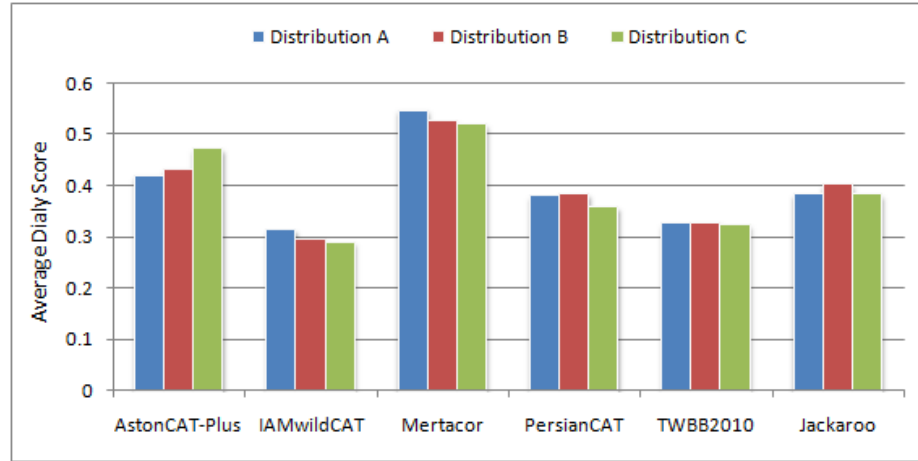


Figure 4.16: Average daily score of each specialist in Type I games.

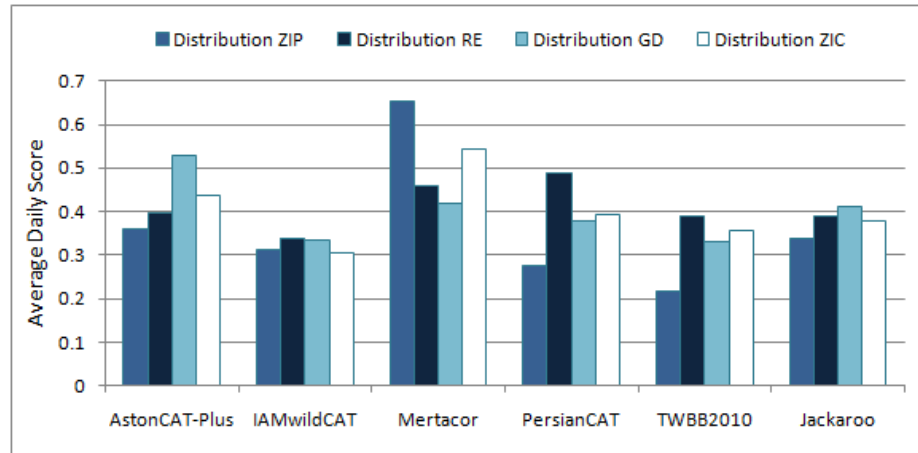


Figure 4.17: Average daily score of each specialist in Type II games.

and are consistent with the rankings of our standard heterogeneous games.

However, in type II games, the trader structure has extreme biases and thus the performance of specialists changed radically. Figure 4.17 shows the scores of type II games, from which we can see Mertacor, AstonCAT and PersianCAT strongly prefer ZIP, GD and RE traders, respectively. IAMwildCAT and Jackaroo have not shown any particular preference for any trading strategy because their overall score variations across different game profiles are far smaller than those of the others.

In the games of type II, Mertacor is no longer the all-time winner. Instead, it ranks the 2nd in Distributions *RE* and *GD*, beaten by PersianCAT and AstonCAT-Plus, respectively. It means that none of the current market mechanism in our controlled experiments is opti-

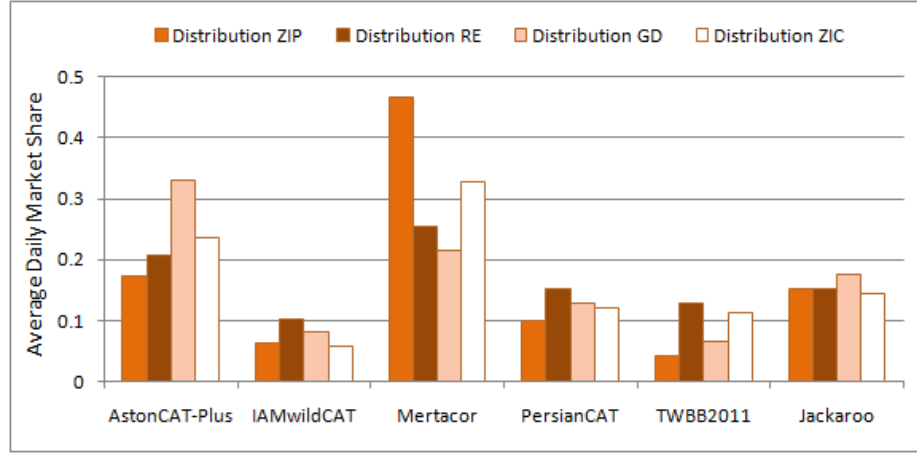


Figure 4.18: Average market share of each specialist in Type II games.

mal. Mertacor's dominance in the ZIP and ZI market and AstonCAT-Plus' dominance in GD market are significant because they respectively score 80.9%, 24.6% and 26.2% higher than the next best entry. In contrast, persianCAT's dominance in the RE market is not very significant since its winning margin is only 6.9%.

Moreover, varying trading strategy distributions do not change the dominance of Mertacor and AstonCAT-Plus in terms of their relative market share. Figure 4.18 shows the dominance of market share by Mertacor and AstonCAT-Plus, because except for RE market (46.3%), the total market share of these two entries is consistently over 50% (63.9% in ZIP, 54.4% in GD and 56.5% in ZI).

In summary, the change of trading strategy mixture definitely affects the score of a specialist. A specialist can show outstanding performance with its preferred trading strategy. Consequently, there is not a specialist that is absolutely optimal against all trader profiles. Nevertheless, a successful specialists show the ability of maintaining their appeal to traders regardless of the dominant trading strategy.

4.4 Conclusion

Firstly, this paper presents a smart e-market agent, called AstonCAT-Plus, which is a post-tournament version of that in the CAT 2010. Specifically, we introduce an effective method to estimate the market equilibrium price by making a good balance between long-term and short-term price tendencies. We also propose some novel market management strategies: (i) the adaptive accepting thresholds for shouts, which depends on dynamic

transaction; (ii) the TPT-CDA clearing strategy, which allows the shout engine to search for more profitable bid-ask pairs to match and clear; and (iii) the hierarchical charging strategy, which can balance well the high market share and the high profit making for the specialist.

Then through a variety of experiments, we evaluate the e-market performance against AstonCAT (tournament version for CAT-2010) and other top entrants of the 2010 competition. The experimental results show that AstonCAT-Plus performs efficiently and stably in heterogeneous games. In particular, it has advantages over other market specialists in terms of *transaction success rate*, *allocative efficiency* and *average trader profit*. It outperforms the original AstonCAT significantly (by 140%) in head-to-head games. Moreover, through experimental analysis we not only demonstrate the strength of AstonCAT-Plus in terms of attracting intra-marginal traders but also identify some features of successful design of double auction market, such as balanced trader profile and high, stable proportion of intra-marginal traders. At last, through the experiment of trading strategy preference games, we discover that no specialist is universally optimal if the traders' strategy distribution is highly biased. However, a successful specialist does show its outstanding ability of maintaining market share across a series of extreme trading strategy compositions.

In the end, we believe both our ideas and findings for market mechanism design are potentially useful to the design and implementation of autonomous e-market of double auction market in the real world. However, we still need to further improve our shout engine method so that the clearing decision can be made on each individual bid-ask pair rather than the matched shouts bunch. Furthermore, our current specialist agent involves too many hard-coded parameters that are manually chosen through experiments. So, it is worth employing an evolutionary approach to learn optimal values for these parameters.

As far as additional experiments are concerned, we will involve the most successful agents from the 2011 CAT tournaments in all kinds of games. To thoroughly explore the effectiveness of TPT-CDA clearing strategy, we will make further comparisons between TPT-CDA and CH-CDA with the same switching points. With similar baselines, we will compare the other individual strategies including strategies from 2011 winner if possible to determine the impact of their strategy on the success of AstonCAT-Plus. Moreover, specific experiments will be designed to find out the amount of contribution of each hierarchical level into the adjustment of fee.

Last but not the least, it is worth developing more realistic autonomous e-markets of double auction, in which the traders could change roles, have bi-directional strategies and dynamic private reservation prices.

Chapter 5

Bi-directional Double Auction and Kernel Trading Strategy

It is well-known that the dominant application of double auction institution is the financial market, in which traders are usually sellers and buyers simultaneously. Hence we introduce a Bi-directional Double Auction (BDA) model, in which the trading activity of every individual trader can be bi-directional. By further introducing a news system to enable traders to update their private valuations, we complete a dynamical financial market simulation.

In BDA market, the selection of trading direction is through trading direction algorithms. Once the decision is made, the order price is determined by a trader's trading strategy. Besides implementing some of the most popular double auction trading strategies (see Section 5.1.5), we develop a new trading strategy called *Kernel* based on probability density estimations, which significantly outperforms all other strategies in our experiments.

The rest of this chapter is organised as follows. Section 5.1 describes the mechanism of BDA market and details each component. Section 5.2 focuses on the details of *Kernel* trading strategy. Section 5.3 shows experimental results and explains our findings. Finally, Section 5.4 concludes the chapter.

5.1 Market Mechanism

This section presents the design of our BDA market.

5.1.1 BDA Market Architecture

Figure 5.1 illustrates the architecture of BDA market. It is composed of three components:

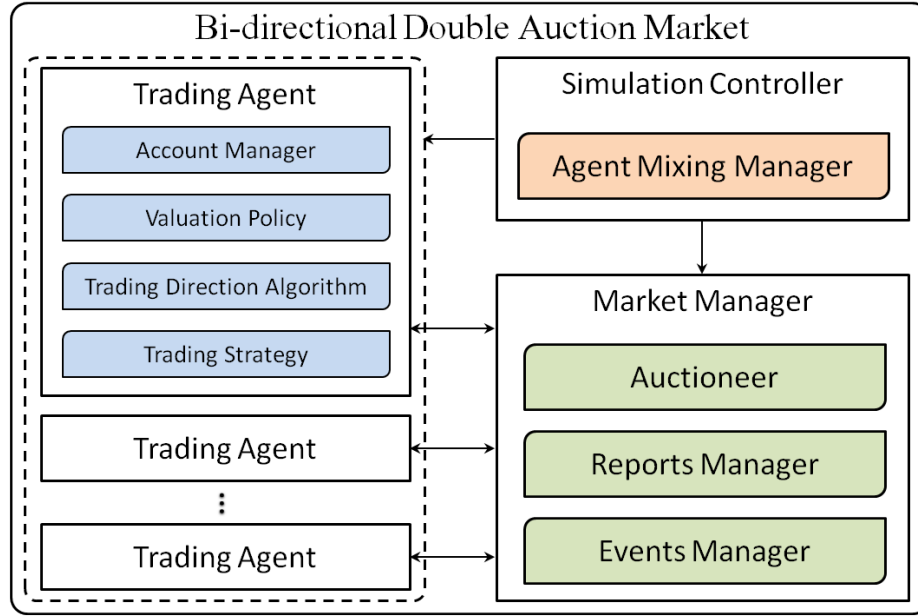


Figure 5.1: Architecture of BDA market.

- Simulation controller** It manages the life cycle (round and day) of the market simulation, initialises the market and traders at the beginning of each game. It contains a class called agent mixing manager, which decides the number of active traders at each round and the way they interact.
- Market manager** It provides all services that market management needs with the help of three following components.
 - (i) **Auctioneer**: It manages orders, generates quotes, clears the market, and so on.
 - (ii) **Reports Manager**: It generates all kinds of reports such as order flow, transaction, equilibrium and group status reports. Traders usually retrieve information they need from various reports to assist their trading decisions.
 - (iii) **Events manager**: It is responsible to broadcast all kinds of events including market events (*e.g.*, order filled and transaction executed events) and simulation events (*e.g.*, round closed and news occurrence events).
- Trading Agent** It is a generic interface for developing traders in BDA market. Trading agents can perceive the status of the market and environment and act accordingly. A trading agent consists of four components:
 - (i) **Account manager**: It manages cash, stock, and daily entitlement accounts. Cash and stock accounts are non-negative such that there is no over-drafting. Once daily entitlement is filled in a day, the trader stops trading until the next day.
 - (ii) **Valuation policy**: It determines how a trader's valu-

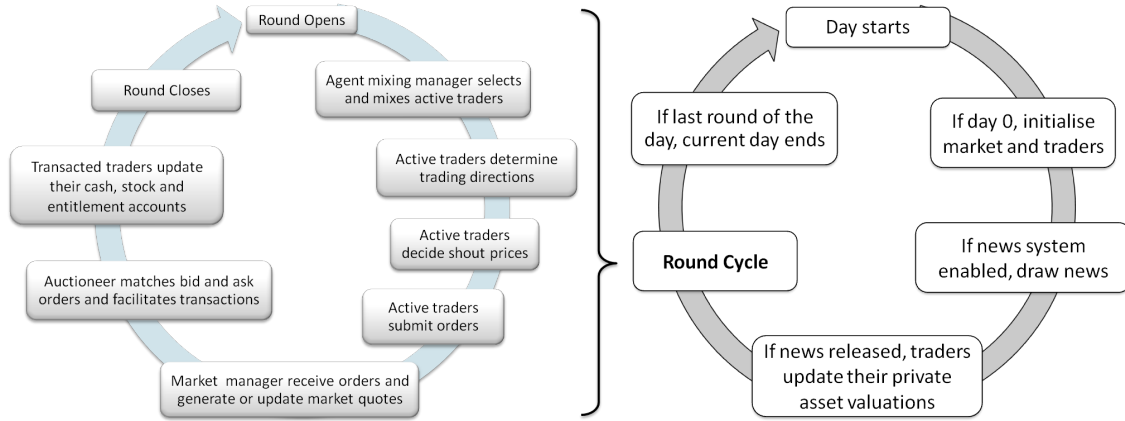


Figure 5.2: Daily activity flow of BDA market.

ation of the asset is initialised and updated. (iii) **Trading direction algorithm:** It generates a trading direction out of *buy*, *sell* and *hold*. (iv) **Trading strategy:** It calculates a shout price should an order be submitted. Trading strategies designed for one-way traders can be employed here.

In terms of implementation, BDA market is implemented as an extension of the Java Auction Simulator API (JASA)¹, which is a high-performance auction simulator designed for performing experiments in agent-based computational economics.

The daily activity flow of BDA market is demonstrated by Figure 5.2. There is a pre-defined population of traders in the market. At the beginning of a day, if news system is enabled, news is drawn with some kind of news occurring probability. If news occurred, traders will update their private valuations accordingly. At the start of a round, agent mixing manager randomly selects a percentage of the traders into the market. In stochastic order, active traders submit new orders or update their previous orders. Each trader will first choose a trading direction based on the decision of his trading direction algorithm. Then, should a trader's decision is to buy or sell one unit of stock, his order pricing strategy will calculate an order price. Afterwards, he submits the order and waits for response from the market. Market manager receives orders and places orders beating current quote on its order book and update market quote. As soon as a match is found, the market is cleared immediately and related events are announced. Traders are informed of these events and update account information as applied. Unmatched orders stay on the order book until order book is reset when resetting conditions are satisfied. Unmatched orders stay alive

¹<http://sourceforge.net/projects/jasa/>, by Steve Phelps

until the end of the day. If the same trader enters the market again at another round and his previous order still exists, his new order will replace the old one.

5.1.2 Agent Mixing Manager

In a real financial market, not every trader is actively trading in every round although the total trading population is relatively stable. Agent mixing manager is introduced to ensure asynchronous order submission in this situation. It selects and mixes active traders in each trading round. Our agent mixing manager uses a *Normal-Poisson* process to determine the number of active traders in each round. This is because Poisson distribution expresses the probability of a given number of events occurring in a fixed interval of time if these events occur with a known average rate [69]. In our case, the event is that a trader enters the market and the average rate is drawn from a *Normal* distribution.

Let $\rho \sim \mathcal{N}(\mu_1, \sigma_1)$ represent the distribution of the average active percentage of the population in each round. A different ρ is drawn at the beginning of every trading day such that the average number of active traders of a round is ρN where N is the number of trader population. The actual number of active traders in a round is $N_a = \text{Pois}(\rho N)$. Finally, N_a active traders are randomly selected from the total population.

5.1.3 Trading Direction Algorithms

Trading direction algorithm is the essential component of a BDA market. Stochastic decision [147] is obviously the simplest trading direction algorithm. But our focus is on the incentive-compatible trading direction algorithms, for which we have developed *Dual* and *Bi*. *Dual* mimics the way normal human traders decide their trading directions in stock markets and therefore is intuitive. *Dual* is simple, fast and not resource-consuming, while generating fairly high allocative efficiency (93.6%). In contrast, *Bi* is more complicated and resource-demanding. However, it is non-parametric and features learning ability based on transaction history. Hence, it achieves higher efficiency in terms of resource allocation (96.1%). In following, we describe *Dual* and *Bi* in details, respectively.

Dual

Dual trading direction algorithm is designed to mimic the way human traders think and behave in choosing a trading direction. The basic assumption is: low(high)-valuation

traders are more likely to submit sell (buy) orders. The closer the market price is to one's private valuation, the more likely that trader will hold.

In *Dual*, traders make decisions by comparing their own valuations with the asset's market prices. Let v be the private valuation of a trader and v_p be the current market price of the traded asset, we introduce α to represent the uncertainty degree (unconfident level) of the trader's private valuation. α takes a value from $(0, 1)$ so that we obtain a valuation spread $[v(1 - \alpha), v(1 + \alpha)]$ for each trader. The trading direction decision is denoted by boolean variables $isBuy$ and $isHold$. When $isHold = true$, irrespective of the value of $isBuy$, trader's decision is to *hold*. When $isHold = false$, the resulting decision is to *buy* if $isBuy = true$ or to *sell* if $isBuy = false$.

Based on whether market price lies inside a trader's valuation spread, the trading direction problem can be divided into two cases: *Deterministic* case if $v_p \notin [v(1 - \alpha), v(1 + \alpha)]$ and *Probabilistic* case if $v_p \in [v(1 - \alpha), v(1 + \alpha)]$. In deterministic case, traders do not *hold*, so

$$isHold = \begin{cases} false & \text{if } v_p < v(1 - \alpha) \\ false & \text{if } v_p > v(1 + \alpha) \end{cases} \quad (5.1)$$

and the decision of buying or selling is quite straight forward:

$$isBuy = \begin{cases} true & \text{if } v_p < v(1 - \alpha) \\ false & \text{if } v_p > v(1 + \alpha) \end{cases} \quad (5.2)$$

In probabilistic case, traders are not sure about their position because their valuations are not definitely higher or lower than the market price. So they are likely to hold. At the same time, because of willing to make a profit, intuitively, the farther v_p is to the left of v on the x-coordinate, the trader is more likely to make a buying decision and vice versa. Since traders' decision are probabilistic, we need to find a way to translate the distance between v and v_p to a probabilistic value rationally. After some research in mathematics, we found sigmoid function very suitable to translate $v - v_p$ into a value between 0 and 1, which can be used to denote the probability of *buy* $P(isBuy)$,

$$P(isBuy) = \frac{1}{1 + e^{-\beta \cdot \lambda (v - v_p)}} \quad (5.3)$$

where $\beta > 0$ is introduced as the trader's risk attitude and λ is a normalization factor. The probability of *sell* $P(isSell)$ is simply $1 - P(isBuy)$. $\beta = 1$ indicates risk-neutral, $\beta < 1$ indicates risk-averse because the sigmoid curve is stretched and the trader is more likely to come up with *hold* decision; and $\beta > 1$ indicates risk-seeking because the sigmoid curve is squeezed and the trader is more likely to come up with *sell* or *buy* decision.

Because every trader's v and α are different, without normalization, calculated probability can easily approach the extreme values of 1 or 0 if $|v - v_p|$ is large enough. If so, the subsequent decision becomes over-deterministic. Hence, the parameter λ in Formula 5.3 is to make sure $P(isBuy) = 0.99$ if and only if $v - v_p = v\alpha$, where $v\alpha$ is the maximum distance between v and v_p in probabilistic case. So λ is derived as follows:

$$\begin{aligned} \frac{1}{1 + e^{-\lambda v\alpha}} &= 0.99 \\ e^{-\lambda v\alpha} &= \frac{1 - 0.99}{0.99} \\ \lambda &= -\frac{\ln \frac{1-0.99}{0.99}}{v\alpha} \end{aligned} \quad (5.4)$$

Due to the symmetric characteristic of sigmoid function, when $v - v_p = -v\alpha$, $P(isBuy) = 1 - 0.99 = 0.01$ and the probability of *sell* reaches the maximum.

Hold is possible in probabilistic case and its probability is calculated by:

$$P(isHold) = \begin{cases} \frac{1-P(isBuy)}{0.5} & \text{if } P(isBuy) > 0.5 \\ \frac{P(isBuy)}{0.5} & \text{if } P(isBuy) \leq 0.5 \end{cases} \quad (5.5)$$

Thus, the closer v_p is to v , the more likely to hold.

Bi

A bid (ask) from a low (high) valuation trader should have a smaller chance of transaction than that from a high (low) valuation trader as long as the offer is “sensible”². Based on this idea, we design trading direction algorithm *Bi* in addition to *Dual*. Because transacted ask or bid prices should appear in different price ranges with different frequencies and these different frequencies can be converted to a probability density curve using kernel technologies, then transaction possibility of future shouts can be estimated.

Since the information about transacted shouts is available in market reports, we can calculate how likely a new shout at the price of a trader's private valuation v is going to be transacted by building kernel density estimators on transacted shout prices. Since a shout can be either a bid or an ask, there are two estimators to be built: bid estimator and ask estimator. After each transaction, two kernel density functions $\mathcal{K}_a(x)$ and $\mathcal{K}_b(x)$ are generated based on the last maximum m transacted bids and asks up to the ones of the last transaction, respectively. Figure 5.3 demonstrates how kernel estimators are built.

²An offer is sensible if the bid price is not greater than valuation or the ask price is not less than valuation

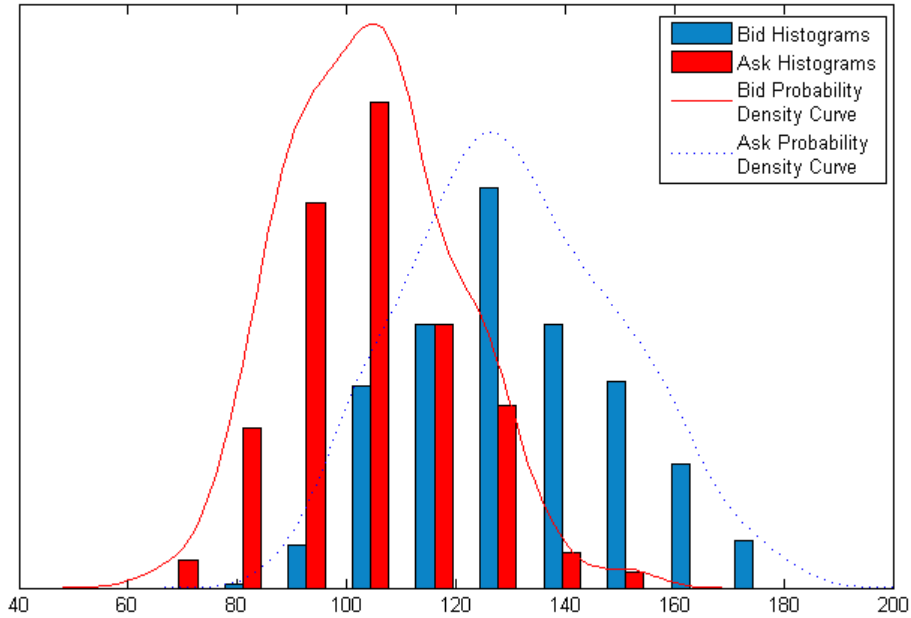


Figure 5.3: Building kernel density estimators.

Blue or lighter (Red or darker) histograms show the frequency that transacted bids (asks) appearing in each bin. Using kernel density estimation techniques, we can build continuous probability density functions for transacted asks and bids respectively, which are shown as red (solid) and blue (dotted) curves respectively in the figure.

Based on $\mathcal{K}_a(x)$ and $\mathcal{K}_b(x)$, we can compute two cumulative probabilities as follows:

$$P_b(v) = \int_{-\infty}^v \mathcal{K}_b(x) dx \quad (5.6)$$

$$P_a(v) = \int_v^{\infty} \mathcal{K}_a(x) dx \quad (5.7)$$

where $P_b(v)$ is the probability of transaction should bid price be v , and $P_a(v)$ is the probability of transaction should ask price be v . Depending on the approximation level of $P_b(v)$ and $P_a(v)$, the probability of *hold* is,

$$P(isHold) = \max(1 - \frac{|P_b(v) - P_a(v)|}{\alpha}, 0) \quad (5.8)$$

In case $isHold = false$, we continue to select a trading direction of *buy* or *sell* whichever got higher transaction possibility, i.e.,

$$isBuy = \begin{cases} true & \text{if } P_b(v) > P_a(v) \\ false & \text{if } P_a(v) > P_b(v) \end{cases} \quad (5.9)$$

5.1.4 News System

The occurrence of news represents the change of environment in BDA market. News is modelled by impact level $\theta \in [-10, 10]$, generated by *simulation controller* with a daily news occurrence probability and released by *events manager*. The news system defines how news happens and influences traders' valuations. The introduction of news system converts a static BDA market into a dynamic one. A positive value indicates bullish news that could cause asset value to rise and vice versa. Once news is received, the most important question is: how should traders' valuations be updated? We use exponential function to model the valuation updating process for its ideal function curve. Let a trader's new valuation be v_n , old valuation be v_o , when one receives a news of impact level θ , his valuation is updated as:

$$v_n = v_o e^{\gamma^2 \theta} + v_p \epsilon \quad (5.10)$$

where $\gamma \in [0.05, 0.20]$ represents the trader's intrinsic news sensitivity. The range of γ is based on the advice of the stock market experts and our experimental results. For the same impact of news, a trader who is more sensitive updates his valuation more aggressively and vice versa.

Nevertheless, $v_p \epsilon$ (where $\epsilon \sim \mathcal{U}(-0.02, 0.02)$) is introduced into the valuation updating function for two purposes: (i) It ensures that two traders with identical γ and private valuation do not get exact the same valuation after updates. (ii) It models the judgement diversion on weak news (*i.e.*, news of low impact levels). The stronger the news, be it positive or negative, the more consistent traders' reactions are because the first part of the formula is dominant over $v_p \epsilon$. The weaker the news, the more diverse trader's reactions become as $v_p \epsilon$ is dominant. Figure 5.4 depicts how a trader's valuation changes excluding $v_p \epsilon$.

5.1.5 Trading Strategies

Most trading strategies used in our BDA market are inherited from well-studies double auction trading strategies including Zero-Intelligence Constrained (ZIC) strategy [63], Zero Intelligence Plus (ZIP) strategy [33], Roth-Erev (RE) strategy [53], Gjerstad-Dickhaut (GD) trading strategy [62]. Besides, we design new trading strategies: *Kernel* and *Forecast*. *Kernel* is detailed in Section 5.2. *Forecast* submits a(an) bid (ask) price that is above (below) a linearly predicted future asset price by a small random amount as well as below (above) the trader's private valuation. In addition, the truth telling strategy (*Truth*) and the stochastic strategy *Noise* are coded for experimentation.

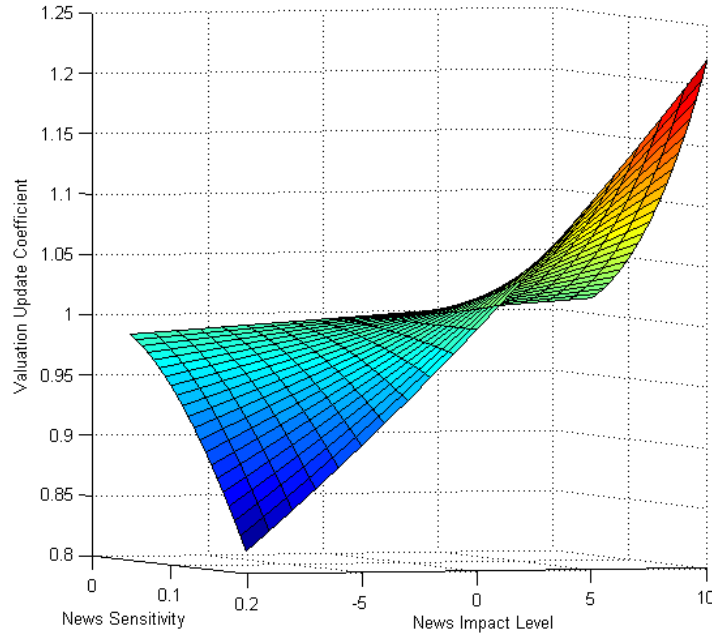


Figure 5.4: Mesh of private valuation update ratio function eliminating the stochastic part.

Details of each implemented trading strategy are given as follows.

1. **ZIC** is Gode and Sunder's Zero-Intelligence Constraint trading strategy [63]. *ZIC* traders determine order prices by adding a uniformly distributed random mark-up to their private valuations. *ZIC* is non-history-based and non-reactive as it does not consider the market condition in its decision-making process [22]. *ZIC* is used as our benchmark trading strategy for its ability of mimicking the trading behaviours of ordinary people from macroscopic view.
2. **ZIP** stands for Zero Intelligence Plus and is first designed by [35] and has been used as a benchmark for strategy evaluation (*e.g.* [43, 166, 179]) in a number of works. In practice, *ZIP* is believed to be frequently used by some mutual funds [66, 67, 80] despite its simplicity relatively. While *ZIC* uses a random mark-up, *ZIP* learns an appropriate profit margin using momentum based algorithms. In BDA market, two learners are separately built to determine selling margins and buying margins respectively.
3. **RE** stands for Roth-Erev algorithm, which is a strategy designed to mimic human game-playing behaviour in extensive form games [53]. This strategy trains a discrete

learner based on immediately feedback from placing previous shouts using 1-armed bandit learning algorithm. Then discrete learner helps the agent to choose an action that balances exploitation of historically profitable actions and exploration of new actions.

4. **GD** trading strategy is developed by Gjerstad and Dickhaut [62]. In *GD*, certain length of shout placement and transaction history are recorded. Then agents use cubic spline interpolation to build a belief function on each successive pair of data item to indicate whether a particular shout is likely to be accepted. Given this information, the bidding strategy is to submit a shout that maximises the trader's own expected surplus.
5. **Forecast** trading strategy borrows the idea of technical analysis which is believed to be widely used in various markets [125, 163, 85]. It uses linear regression to predict a future asset price p_f based on selective historical prices. Because the agent believes p_f will be the next transaction price, his bid and ask prices are $p_f(1 + \epsilon_f)$ and $p_f(1 - \epsilon_f)$ where $\epsilon_f \sim \mathcal{U}(0, 0.06)$, respectively to suppress the transaction cost while standing reasonable chance of obtaining the transaction opportunity. There is never lack of criticism from some sectors of the industry and academia on technical analysis. But there are also numerous works suggesting its profitability [126, 19, 5, 76, 47, 49, 48]. As an evidence, *Forecast* strategy has shown very good (better than *GD*) performance in many of our experiments.
6. **Kernel** trading strategy determines shout prices based on a kernel probability density estimator built on historical transaction information to calculate the possibility of a future shout being transacted. The details are given in Section 5.2.
7. **Truth** strategy simply bids the trader's private valuation on the established trading direction.
8. **Noise** strategy chooses trading direction randomly. Then on the selected trading direction, it forms a shout prices using *ZIC* trading strategy.

5.2 Kernel Trading Strategy

Kernel trading strategy is constructed based on kernel probability density estimators built upon historical transaction data in BDA market, which calculate the chance of a shout

(ask or bid) being transacted. It has close relationship to *Bi* because they use the same kernel estimators. After trading direction is generated (by *Dual* or *Bi*), kernel estimators (updated automatically after every transaction) are employed to assess the profitability of each possible shout on the chosen trading direction. Finally, the most profitable shout price will be selected for submission.

Assuming in the last m transactions, the lowest transacted bid price is \underline{b} and the highest transacted ask price is \bar{a} . Thus, the searching spaces for the optimal bid and ask are $[\min(0, \underline{b}(1 - 0.05) - 0.05v), v]$ and $[v, \bar{a}(1 + 0.05) + 0.05v]$, respectively. As seen, we allow both relative (5%) and absolute ($0.05v$) extra space to explore the best possible price outside transacted bids or asks. Meanwhile, the extra space is limited to keep the algorithm efficient. Moreover, we use $\mathcal{K}'(p)$ to denote the transaction probability of the price point p on the estimated probability density curve. Thus, the optimal bid b^* or optimal ask a^* can be calculated as follows:

$$b^* = \arg \max_{p \in [\min(0, \underline{b}(1 - 0.05) - 0.05v), v]} \mathcal{K}'_b(p) \cdot (v - p) \quad (5.11)$$

$$a^* = \arg \max_{p \in [v, \bar{a}(1 + 0.05) + 0.05v]} \mathcal{K}'_a(p) \cdot (p - v) \quad (5.12)$$

Because $\mathcal{K}(x)$ is a probability density function, $\mathcal{K}(p)$ is actually represented by the probability of a small area around p . The small area is defined by a precision threshold δ . So, $\mathcal{K}'(p)$ can be obtained by,

$$\mathcal{K}'(p) = \int_{p-\delta}^{p+\delta} \mathcal{K}(x) dx \quad (5.13)$$

where the default value of δ is 0.01.

5.3 Controlled Experiments

The framework of our BDA market is modular and highly customisable. By setting parameters differently, we can easily simulate many different market scenarios. For example, it can be set to continuous BDA market or Clearing House (CH) BDA market. By default, all experiment results discussed in this thesis are based on continuous BDA market. Basically, we have run two types of experiments: (i) static games where *hold*, news system and agent mixing manager are disabled and (ii) dynamic games where all disabled dynamic features are enabled. By default, all agents use *Dual* trading direction algorithm. Games with random initial valuations are run for 100 iterations. Games with fixed-array initial valuations are run for 20 iterations. In each iteration, there are 100 days of 10 rounds daily

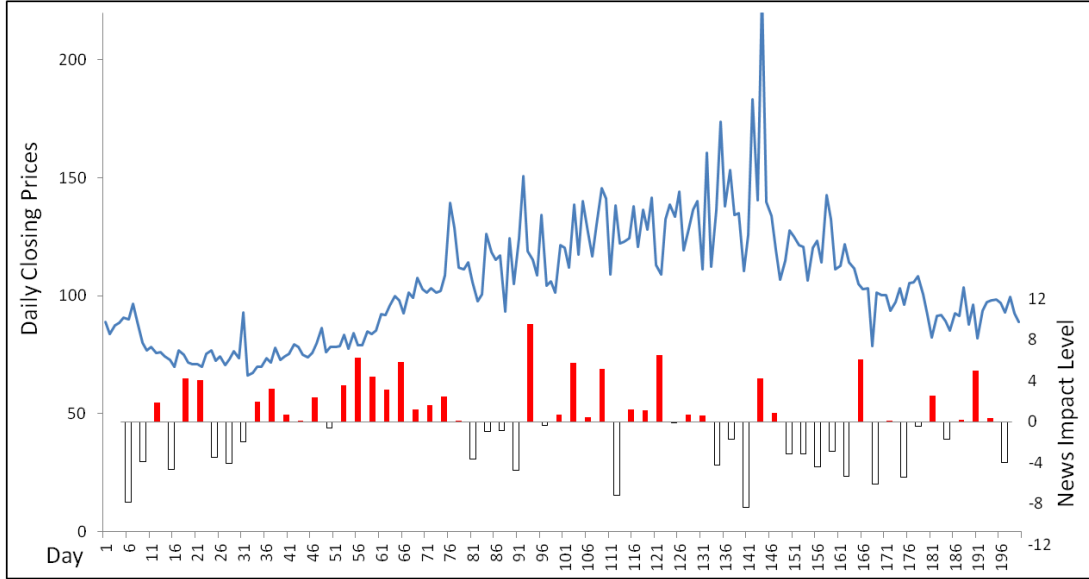


Figure 5.5: Daily closing prices in iteration 6 and corresponding news series on the same time scale.

News impact levels are drawn from Normal distribution $\mathcal{N}(0, 4)$. The probability of news occurrence on a day is 0.3.)

for static games or 200 days of 50 rounds daily for dynamic games. In the performance comparisons, usually the aggregate results of trader groups covering the same valuation distribution or valuation array are compared. Initial private valuations are drawn from a uniform distribution $\mathcal{U}(60, 120)$ in static games, and assigned by a fixed-array in dynamic games. Table 5.1 shows the details of the configurations of our experiments.

5.3.1 BDA Market Time Series

In dynamic games, BDA market produces rational time-series in response to news. Figure 5.5 shows how daily closing price evolves in response to news impact levels during 200 simulation days (data taken from iteration 6). The asset value generally rises with positive news and falls with negative ones, which agrees with our intuitions.

In terms of return, BDA market can reproduce some well-known stylised facts of real financial markets. The distribution of daily returns displays a heavy tail with positive excess kurtosis around 2.48. The absolute daily log returns show a decaying auto-correlation as time lags increase (see Figure 5.6).

When the news system is disabled, asset value do not shift in general because the shape

Features/Game	Efficiency Game 1	Efficiency Game 2	Efficiency Game 3	Efficiency Game 4	Static Heterogeneous Game	Dynamic Heterogeneous Game
Game type	Static	Static	Static	Static	Static	Dynamic
News system	Disabled	Disabled	Disabled	Disabled	Disabled	Enabled
Is HOLD available	No	No	No	No	No	Yes
News occurring probability	N/A	N/A	N/A	N/A	N/A	0.3
News impact distribution	N/A	N/A	N/A	N/A	N/A	Normal(0,4)
Number of days	100	100	100	100	100	200
Num of rounds/Day	10	10	10	10	10	50
Total trader population	100	100	100	100	240	240
Active percentage in each round	100%	100%	100%	100%	100%	30%
Active Percentage perturbation	0%	0%	0%	0%	0%	5%
Normal-Poisson Agent mixing	Disabled	Disabled	Disabled	Disabled	Disabled	Enabled
Numer of groups	1	1	1	2	8	8
Number of trader in each group	100	100	100	50	30	30
Grouping criteria	Trading Direction Algorithm	Valuation Fuzzy Level	Trading Strategy	Trading Direction Algorithm	Trading Strategy	Trading Strategy
Trading direction algorithm	Dual Bi Stochastic	Dual	Dual	Dual Bi	Dual	Dual
Trading strategy	ZIC	ZIC	Varies	ZIC	Varies	Varies
Trader daily entitlements	5	5	5	5	5	10
Valuation policy	Random	Random	Random	Random	Random	Array
Random valuation setting	Uniform(60,120)					
Array valuation setting	{61,64,68,70,72,74,76,78,79,80, 82,83,84,86,88,90,92,93,94,95, 96,98,100,101,104,106,108,110,114,119}					
Iteration	100	100	100	100	100	20

Table 5.1: BDA market experiment configuration details.

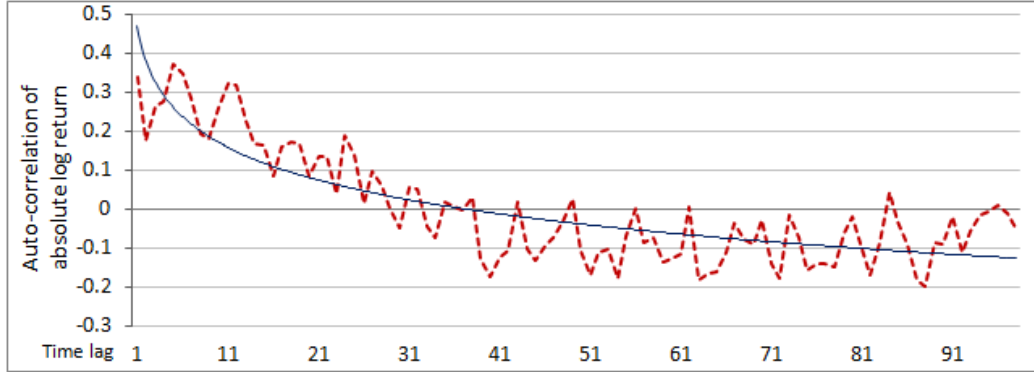


Figure 5.6: Decaying auto-correlation of absolute daily log returns in BDA market.

Direction Algorithm	Mean	Standard Deviation
Bi	0.961	0.0113
Dual	0.936	0.0142
Stochastic	0.699	0.0171

Table 5.2: Market average efficiency and standard deviation calculated based on the data of 100 static game iterations.

of the curve of the expected probability density of transaction prices is perfectly balanced (see Figure 5.7).

5.3.2 Static Market Efficiency

One of the essential research topics about double auction markets is the market's allocative efficiency, *i.e.*, the ratio of actual profit of all traders to the theoretically maximum profit. Therefore, static games are designed to investigate the efficiency of BDA market. ZIC [63] is used as our benchmark trading strategy as it is proven one of the most successful applications of agent-based computational economics [96] and has been found to mimic trader behaviour closely [45]. Our first finding is that trading direction algorithms seriously affect the market efficiency. As shown in Figure 5.8 and Table 5.2, stochastic trading directions lead to poor market efficiency and biggest variance. As long as the trading direction algorithm is incentive compatible, the market efficiency is significantly improved from 69.9% to minimally 93.6%. The intelligent *Bi* beats the intuitive *Dual* by 2.5%. *Bi* also keeps the efficiency above 95% for 87% of time.

The market efficiency is also affected by traders' confidence with their private valuations. The greater the uncertainty, the lower the confidence. In relevant experiments, we set

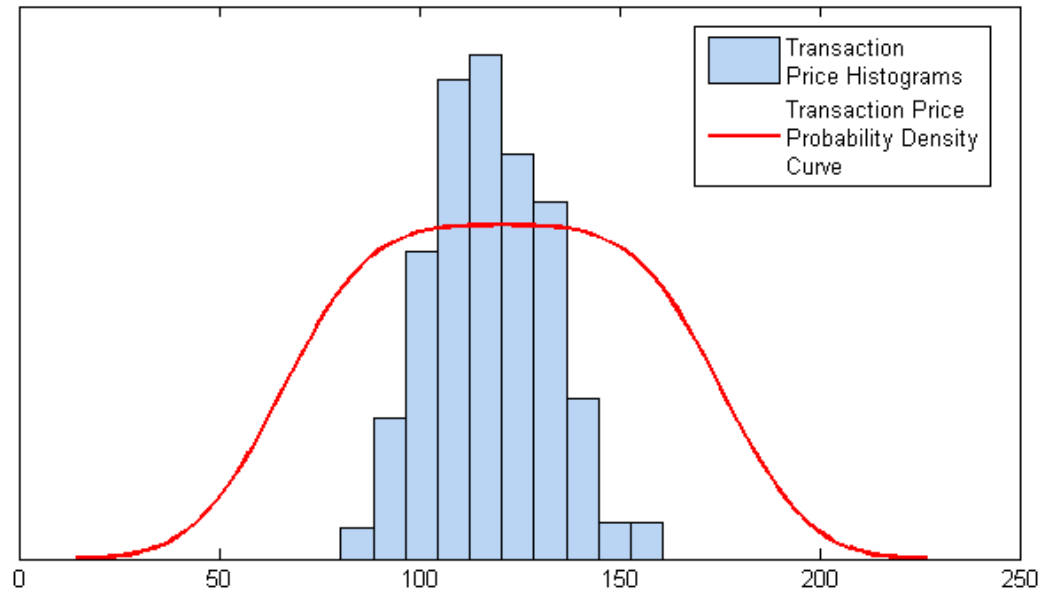


Figure 5.7: Convergence of transaction prices in static games.

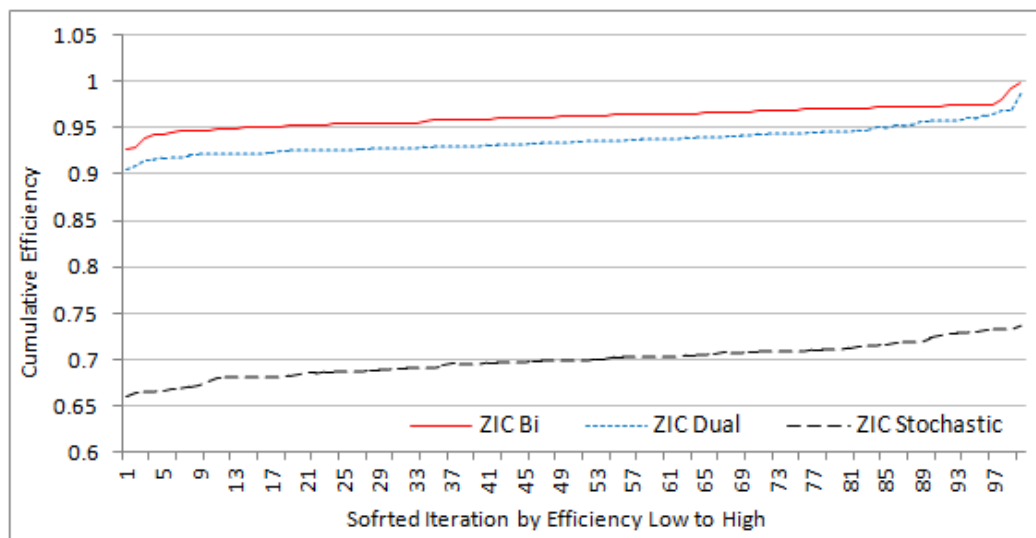


Figure 5.8: Sorted market efficiencies of 100 static games.

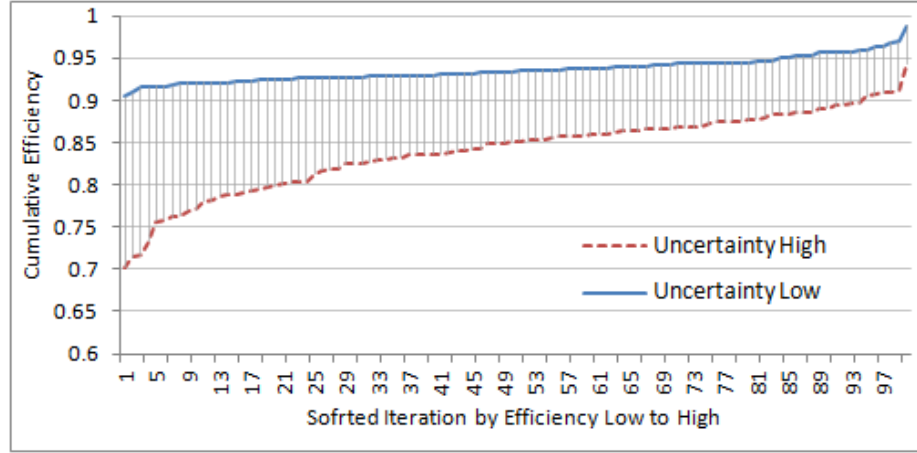


Figure 5.9: Sorted market efficiencies categorised by uncertainty degree distributions of the valuation.

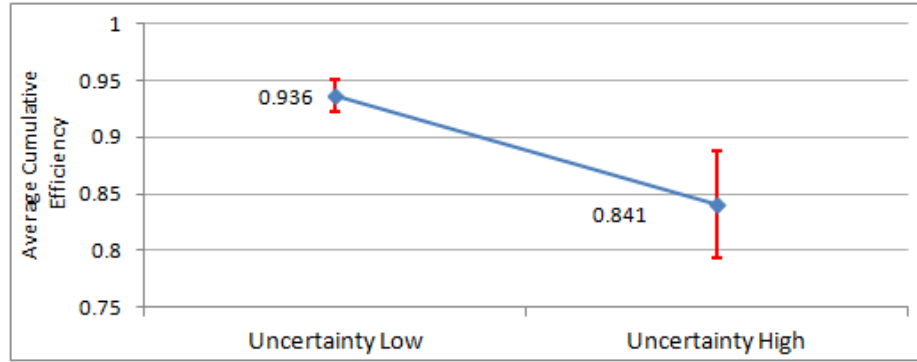


Figure 5.10: Average market efficiencies of static game in regard to uncertainty distributions of the valuation.

$\mathcal{U}(0, 0.15)$ as traders' uncertainty degree distribution for the first batch and $\mathcal{U}(0.15, 0.35)$ for the second batch. Each batch contains 100 iterations and results are shown in Figures 5.9 and 5.10. It can be seen that the high confidence (low uncertainty) market significantly outperforms the low confidence (high uncertainty) market. Additionally, the high confidence market demonstrates better stability as its efficiency deviation is only 0.014 while the low confidence market's deviation is 0.047. The minimum allocative efficiency is significantly improved in the high confidence market compared to the low confidence one (0.905 vs 0.702).

Figure 5.11 shows how the cumulative market efficiency typically evolves during a game. We notice that without a rational trading direction strategy, market efficiency starts

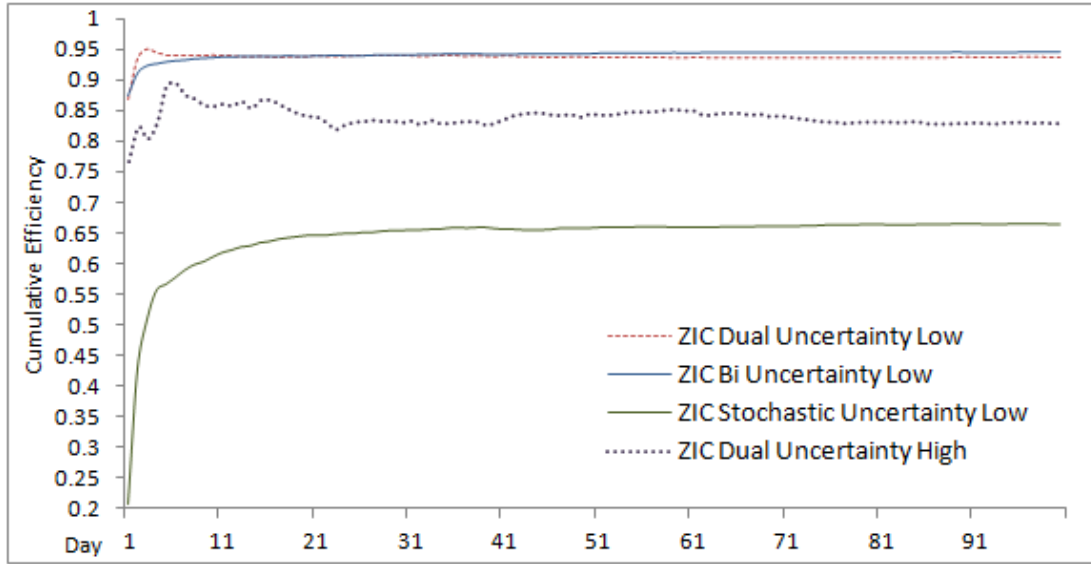


Figure 5.11: Daily cumulative market efficiency in a game.

as low as 0.2 and also finishes significantly lower than those with rational ones. *Bi* shows good learning ability. Its efficiency is not as high as *Dual* when the game starts, but overtakes *Dual* after about 20 days. Uncertainty degree α has a significant influence on the convergence of efficiency as well as market stability. The efficiency of the low uncertainty market is much stabler from day to day, and converges much faster (in 10 days) than the high uncertainty market.

Finally, the market efficiency is also affected by trading strategy (see Figure 5.12). GD and Kernel strategies cannot trade with themselves although they are the best performers in heterogeneous environments. After inspection, we found their calculated optimal buying prices are lower than selling prices. Consequently, data of these two strategies are not displayed in Figure 5.12. In summary, if every trader in a BDA market has the ability to calculate a near-optimal price of bid or ask after sensible choice of trading direction, transaction opportunities among them will be reduced seriously. So we infer, in a financial market where the asset's fundamental value does not change, there must be unintelligent traders to make exchanges happen.

For other strategies, Figure 5.12 shows their average efficiencies and standard deviations. The market reaches the best allocative efficiency when filled with truth telling traders. As long as traders select trading directions rationally, they virtually need to do nothing extra to achieve high allocative efficiencies. However, if traders all want to exploit others using intelligent trading strategies, the more intelligent the trading strategies are, the

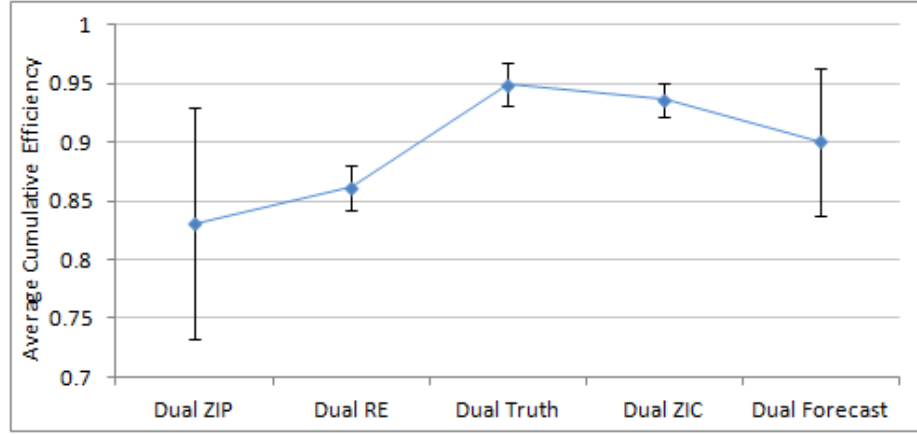


Figure 5.12: Average market efficiencies and deviations regarding trading strategies.

less efficient the market is.

5.3.3 Performance of Kernel Strategy

Our *Kernel* trading strategy shows excellent performance in both static and dynamic games. In static heterogeneous market, traders of multiple types compete to make their own profit. Every group employs *Dual* as trading direction algorithm such that the different group profit is the result of different trading strategies. As shown by Figure 5.13, the highest average group profit comes from kernel group. According to paired 1-tail student t-test performed over results of 100 iterations, the probability that Kernel group's profit has the same mean to the second best group is 3.29×10^{-7} . GD is the strategy that achieves the highest unit profit (see Figure 5.15). However, with the smallest quantity of all (see Figure 5.14), GD's final performance in terms of total profit is only modest. Dual ZIP and Truth traders have traded large quantities at the cost of giving profit away to competitors, which leads to their poor performance in terms of total profit. Therefore, the most profitable strategy should earn good profit in each transaction while not letting transaction opportunities slip away due to its unfavourable shout prices.

We also investigate whether trading direction algorithms have a significant contribution to the profitability. Our answer is positive as *Dual*'s transaction quantity is only 92.9% of *Bi* while *Dual*'s unit profit is also slightly lower than *Bi*. Consequently, *Bi* ZIC traders make about 10% more profit when compete with *Dual* ZIC traders in one market.

In dynamic heterogeneous games, multiple types of traders compete to maintain their wealth. News events are generated exogenously with occurrence probability of 0.3. News

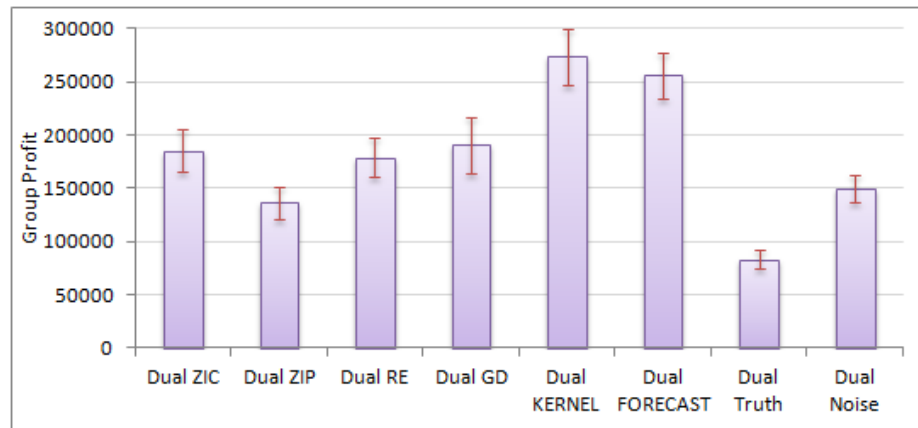


Figure 5.13: Average profit of each trader group in static heterogeneous game (100 iterations).

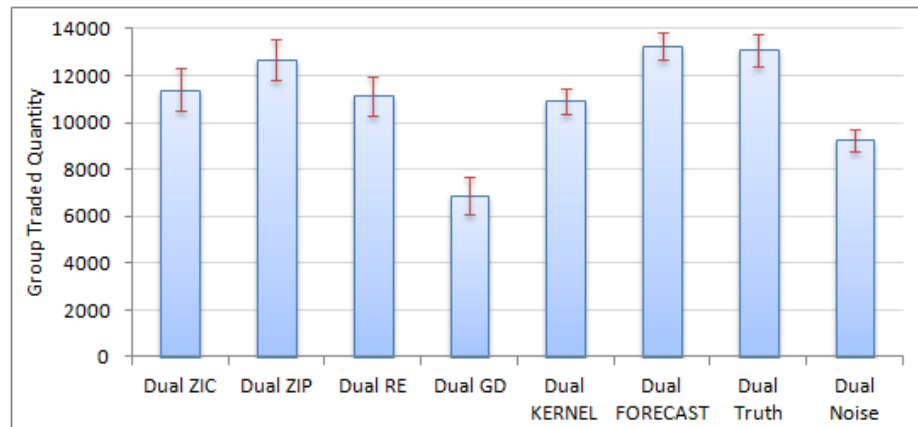


Figure 5.14: Average traded quantity of each trader group in static heterogeneous game (100 iterations).

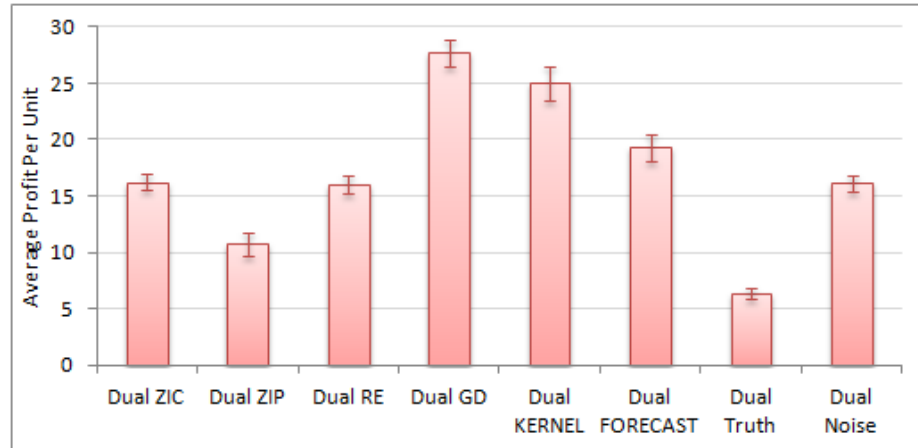


Figure 5.15: Mean of the average profit per unit of each trader group in static heterogeneous game (100 iterations).

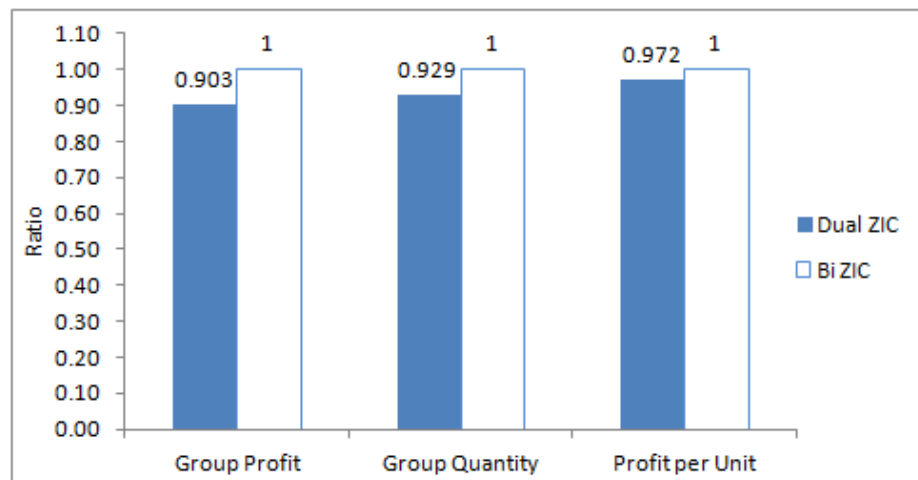


Figure 5.16: Proportional comparison of profit, traded quantity and unit profit between *Dual ZIC* trader group and *Bi ZIC* trader group (100 iterations).

impact levels are subject to Normal distribution $\mathcal{N}(0, 4.0)$. As response to news, traders update their private valuations asynchronously, which cause the Rational Expected Equilibrium (REE) to shift continuously. Every group starts with exactly the same wealth of 11,356,800 and valuations. Their news sensitivities are from the same distribution.

After 20 iterations, *Kernel* traders achieve the best aggregate wealth. From Figure 5.18, we can see kernel group's average wealth exceeds that of GD group by 1.36% and ZIC group by 4.91%. Although the advantage is marginal in terms of ratio, in financial market, even 1% difference could mean millions of pounds of profit or loss. Averaged over all iterations, the winner group's wealth surpass the worst group by 8.4%. Clearly, an effective trading strategy is needed if the investor wants to maintain his wealth properly. Figure 5.17 provides the details of the final wealth of every group in each iteration of dynamical heterogeneous wealth game.³ Kernel group's wealth is ranked first for 18 times out of 20 iterations. The black broken line indicates group initial wealth level. In a falling market, kernel traders maintain their wealth best without exception. In a rising market, GD group wins twice at iterations 13 and 20 and kernel group wins the rest.

Obviously, GD traders perform better in a rising market, which is to do with their low volume (*i.e.*, the smallest traded quantity among all groups). Specifically, GD's average quantity is 20505.3, which is 55% of Kernel's and only 37% of that of Truth group that has the largest average volume. Low volume means GD traders keep more stocks in their own hands in stead of exchanging them with other traders. In a rising market, their stocks rise in value, which leads to their good performance in terms of maintaining wealth. In a declining market, this advantage does not exist any more. The worst performing group is the Truth group. Members of the group are the most frequent traders in the market. However, they only make averagely 8.655 profit in each transaction. Thus, the more they trade, the more loss they make in each transaction. Since everyone's stock is limited, they are actually giving wealth away to their trading counterparts. This corresponds to the phenomena in real market that non-strategic traders who frequently trade but becomes the worst wealth maintainers.

5.4 Conclusion

In this chapter, we present the design of bi-directional double auction market together with the new *Dual* and *Bi* trading direction algorithms to specifically model how trading

³We have also done the same experiments with *Bi* and obtained the same results statistically.

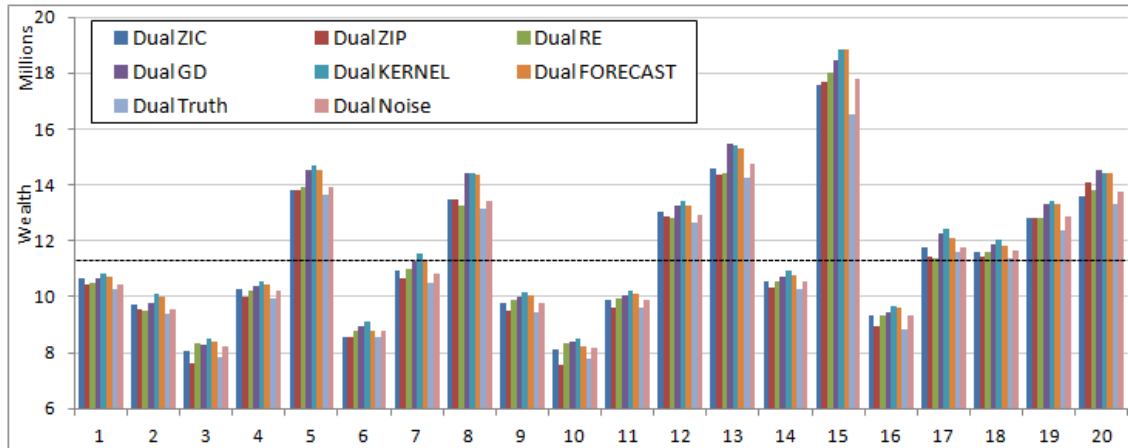


Figure 5.17: Final wealth each trader group possess in dynamic heterogeneous games. From left to right, strategy of the group is Dual ZIC, Dual ZIP, Dual RE, Dual GD, Dual Kernel, Dual Forecast, Dual Truth and Dual Noise (20 iterations).

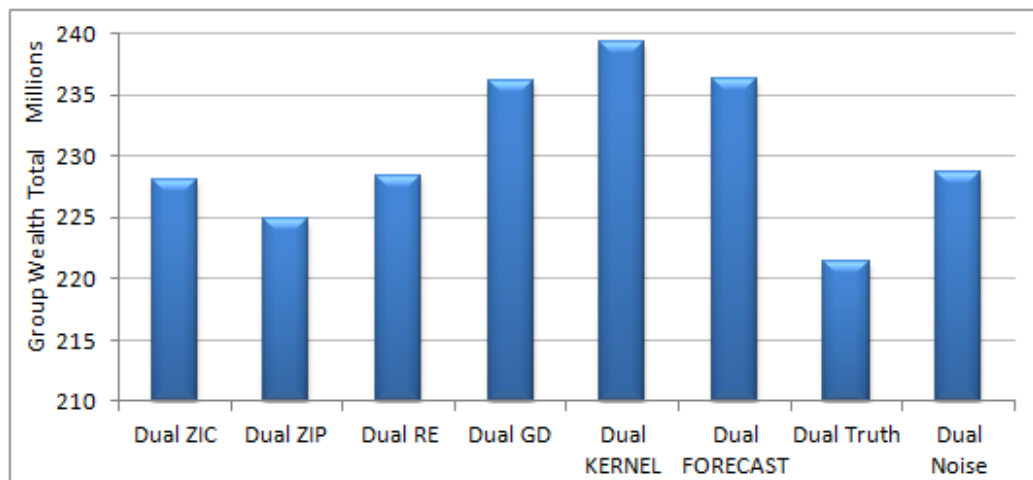


Figure 5.18: Aggregate group wealth of dynamic heterogeneous games (20 iterations).

directions are dynamically decided. In addition, the news system and *Normal-Poisson* agent mixing manager are introduced to complete a dynamic financial market simulation. Through experiments, we find that BDA market generates rational series against varying impact levels of news and successfully reproduces some typical stylised facts of real financial markets. In terms of the allocative efficiency of a static continuous BDA market, the trading direction algorithm is found to be the most critical factor and intelligence of trading strategies has the effects of reducing the market efficiency provided that traders' trading direction decisions are rational. Furthermore, our new *Kernel* trading strategy demonstrates superior performance to others in terms of both making profit in static games and maintaining wealth in dynamic games.

Chapter 6

Conclusion and Future Work

Nowadays, e-commerce is gradually taking the primary position of the traditional business. One can almost purchase any commodity or service on the Internet without going out to the brick-and-mortar shops or meeting any real person. Therefore, we have the reasons to believe that more and more traditional trading via person-to-person negotiation will be replaced by the autonomous interactions among software agents on behalf of their owners in the future. Subsequently, designing efficient economic mechanisms such as the market system and developing effective trading strategies in any specific economic system are two very important and meaningful topics, which are exactly the aims of this work.

The first economic system we investigate is the sponsored search in which GSP-based keyword auctions are the main method to determine prices and resource allocations. AstonTAC is a trading agent designed for such auction. It is shown to be successful and stable across a wide range of TAC AA environments in both the competition and our controlled experiments. In particular, we attribute its success to the bidding price generating strategy and the query selector. Market-based Value Per Click reflects the dynamic change of market value of different keywords and thus leads to the generation of flexible and adaptive bidding prices. In a multi-keyword scenario especially when keywords are interdependent, it is essential to view the set of all possible queries as a whole while conceiving a strategy. Selective bidding is one of the ways. Query selector makes sure that we invest only on the most profitable keywords of each day. Given a limited conversion capacity, knowing how much further we can go beyond it gives us ability of grabbing the largest possible number of conversion to extend our profit space.

Although, strategies employed by AstonTAC are tailored to the specific context of the AA competition, due to the similar features of the TAC AA competition and real sponsored search, we believe that concepts developed for AstonTAC are broadly applicable to an

advertiser agent in the real sponsored search scenario.

Following the research in GSP, we switch our attention to a more influential auction format, namely Double Auction, which is employed to determine transaction prices and resource allocations of many powerful economic systems represented by the financial market. AstonCAT-Plus is the post-tournament version of AstonCAT in the TAC CAT tournament which concentrates on the investigation of double auction market mechanisms. It is shown to be an efficient, stable, profitable and trader-friendly e-market in our controlled experiments. We attribute its success to the following strategies: (i) We introduce an effective method to estimate the market equilibrium price by merging long-term and short-term transaction price information. (ii) We design an adaptive accepting strategy which contributes to its high transaction success rate. (iii) Our TPT-CDA clearing strategy, which allows the shout engine to search for more profitable bid-ask pairs to match and clear, which improves the profit of traders especially intra-marginal traders. (iv) Our hierarchical charging strategy balances the needs of attracting traders and make much as possible profit for the e-market agent.

Through a variety of experiments, we evaluate AstonCAT-Plus's performance against AstonCAT (tournament version for CAT-2010) and other top entrants of the 2010 competition. Our experimental results show that AstonCAT-Plus performs efficiently and stably in heterogeneous games. In particular, it has advantages over other market specialists in terms of *transaction success rate*, *allocative efficiency* and *average trader profit*. It outperforms the original AstonCAT significantly (by 140%) in head-to-head games. Moreover, through empirical analysis, we not only demonstrate the strength of AstonCAT-Plus in terms of attracting intra-marginal traders but also identify some features of successful design of double auction market, such as balanced trader profile and high, stable proportion of intra-marginal traders. At last, in trading strategy preference games, we uncover that no specialist is universally optimal if the traders' strategy distribution is highly biased. However, a successful specialist does show its outstanding ability of maintaining market share across a series of extreme trading strategy compositions. We believe both our ideas and findings for market mechanism design are potentially useful to the design and implementation of automated double auction market in the real world.

In future, we will improve our shout engine method so that the clearing decision can be made on each individual bid-ask pair rather than the matched shouts bunch. To thoroughly explore the effectiveness of TPT-CDA clearing strategy, we will make further comparisons between TPT-CDA and CH-CDA with the same switching points. We will also design specific experiments to find out the amount of contribution of each hierarchical level into

the adjustment of fee. Nevertheless, our current specialist agent involves many hard-coded parameters that are manually chosen through experiments. So, it is worth employing a machine learning approach, such as an evolutionary approach to learn optimal values for these parameters.

CAT competition is developed on the basis of a traditional double auction platform - JASA. By running experiments on such a platform, we realise that traders are assigned a fixed trading direction before the simulation starts. But real financial traders can change their trading direction at any time. If traders can decide their trading direction dynamically, would the market show different properties? With this question, we design and implement the BDA market. By introducing a news system, we do not only try to explore the properties of a static BDA market, but also explain phenomenon emerged in financial markets using our financial market simulation model. Consequently, *Dual* and *Bi* trading direction algorithms are developed to model how trading directions are dynamically decided.

Through experiments conducted in the continuous BDA market, we have the following findings. First, BDA market generates rational series against varying impact levels of news and successfully reproduces some typical stylised facts of real financial markets. The allocative efficiency of a static BDA market largely comes from the rational trading direction selections of the traders. With sensible trading direction algorithms, the intelligence of trading strategy has the effects of reducing the market efficiency. A high-confidence market is more efficient and stable than a low-confidence one. Moreover, we design a high-performance trading strategy called *Kernel* in the BDA market, which also works in DA markets. *Kernel* strategy utilises the techniques of probability density estimation to seek the optimal bidding price after a trading direction is decided. Kernel trading strategy demonstrates superior performance to others in terms of both making profit in static games and maintaining wealth in dynamic games.

In future, we plan to extend this work in the following directions. (i) We will conduct additional experiments to further explore the properties of the BDA market, *e.g.*, to compare different ways of extracting the asset's market price, to explore the influence of risk attitudes, to explore the properties of the BDA market in the setting of Clearing House and to investigate how different new impact distributions affect the performance of each trading strategy. (ii) We will implement and study other well-known trading strategies into the BDA market, such as adaptive-aggressive strategy [176], extended GD strategy [165] and fuzzy-logic based strategies [72]. (iii) We will improve our model according to real financial data and possibly extend it into a distributed system that accepts human offers from remote machines. Then we will be able to find out in the BDA market whether humans can

beat computer programs.

Bibliography

- [1] Z. Abrams and M. Schwarz. Ad auction design and user experience. *Applied Economics Research Bulletin*, 2(1):98–105, 2008.
- [2] A. Agarwal, K. Hosanagar, and M.D. Smith. Location, location, location: An analysis of profitability of position in online advertising markets. *Journal of marketing research*, 48(6):1057–1073, 2011.
- [3] G. Aggarwal, A. Goel, and R. Motwani. Truthful auctions for pricing search keywords. In *Proceedings of the 7th ACM conference on Electronic commerce*, pages 1–7. ACM, 2006.
- [4] B. Alexandrova-Kabadjova, E. Tsang, and A. Krause. Market structure and information in payment card markets. *International Journal of Automation and Computing*, 8(3):364–370, 2011.
- [5] F. Allen and R. Karjalainen. Using genetic algorithms to find technical trading rules. *Journal of financial Economics*, 51(2):245–271, 1999.
- [6] J. Arifovic. Evolutionary dynamics of currency substitution. *Journal of Economic Dynamics and Control*, 25(3):395–417, 2001.
- [7] W.B. Arthur, J.H. Holland, B. LeBaron, R.G. Palmer, and P. Talyer. Asset pricing under endogenous expectations in an artificial stock market. In S. Durlauf W. B. Arthur and D. Lane, editors, *The Economy as an Evolving Complex System II*. Addison-Wesley, 1997.
- [8] R. Axelrod. Simulation in the social sciences. *Handbook of research on nature inspired computing for economy and management*, pages 90–100, 2007.
- [9] E. Basker. Job creation or destruction? Labor market effects of Wal-Mart expansion. *Review of Economics and Statistics*, 87(1):174–183, 2005.
- [10] RG Bates, MAH Dempster, and Y.S. Romahi. Evolutionary reinforcement learning in FX order book and order flow analysis. In *Computational Intelligence for Financial Engineering, 2003. Proceedings. 2003 IEEE International Conference on*, pages 355–362. IEEE, 2003.

- [11] A. Beltratti and S. Margarita. Evolution of trading strategies among heterogeneous artificial economic agents. In *Proceedings of the second international conference on From animals to animats 2: simulation of adaptive behavior: simulation of adaptive behavior*, pages 494–501. MIT Press, 1993.
- [12] A. Blum, T. Sandholm, and M. Zinkevich. Online algorithms for market clearing. *Journal of the ACM (JACM)*, 53(5):845–879, 2006.
- [13] L. Blume, D. Easley, and M. O’hara. Market statistics and technical analysis: The role of volume. *The Journal of Finance*, 49(1):153–181, 2012.
- [14] K. Boehret. A deal on a haircut? That’s what friends are for. *Wall Street Journal*, 2010.
- [15] B. Boer-Sorbàn. *Agent-Based Simulation of Financial Markets - A Modular, Continuous-Time Approach*. PhD thesis, Erasmus School of Economics, Erasmus University Rotterdam, 2008.
- [16] F. Brandt and G. Weiß. Antisocial agents and vickrey auctions. *Intelligent Agents VIII*, pages 335–347, 2002.
- [17] J. Bredin, D.C. Parkes, and Q. Duong. Chain: A dynamic double auction framework for matching patient agents. *Journal of Artificial Intelligence Research*, 30(1):133–179, 2007.
- [18] R. Briggs and N. Hollis. Advertising on the web: is there response before click-through? *Journal of Advertising Research*, 37:33–46, 1997.
- [19] W. Brock, J. Lakonishok, and B. LeBaron. Simple technical trading rules and the stochastic properties of stock returns. *The Journal of Finance*, 47(5):1731–1764, 2012.
- [20] W.A. Brock and C.H. Hommes. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic dynamics and Control*, 22(8-9):1235–1274, 1998.
- [21] J. Bulow and P. Klemperer. Prices and the winner’s curse. *RAND journal of Economics*, pages 1–21, 2002.
- [22] K. Cai, E. Gerding, P. McBurney, J. Niu, S. Parsons, and S. Phelps. Overview of CAT: A market design competition. Technical report, Technical Report ULCS-09-005, Department of Computer Science, University of Liverpool, Liverpool, UK, 2009. Version 2.0, 2009.

- [23] K. Cai, J. Niu, and S. Parsons. On the economic effects of competition between double auction markets. *Agent-Mediated Electronic Commerce and Trading Agent Design and Analysis*, pages 88–102, 2010.
- [24] M. Cary, A. Das, B. Edelman, I. Giotis, K. Heimerl, and A. Karlin. Greedy bidding strategies for keyword auctions. In *Proceedings of the 8th ACM conference on Electronic commerce*, pages 262–271, 2007.
- [25] P.A. Castro and J.S. Sichman. Towards cooperation among competitive trader agents. *Proceedings of 9th ICEIS., Funchal, Portugal*, pages 138–143, 2007.
- [26] P.A. Castro and J.S. Sichman. Agex: A financial market simulation tool for software agents. *Enterprise Information Systems*, pages 704–715, 2009.
- [27] D. Challet, M. Marsili, and Y.C. Zhang. Modeling market mechanism with minority game. *Physica A: Statistical Mechanics and its Applications*, 276(1):284–315, 2000.
- [28] M. Chang, M. He, and X. Luo. Designing a successful adaptive agent for tac ad auction. In *Proceedings of the 2010 conference on ECAI 2010: 19th European Conference on Artificial Intelligence*, pages 587–592. IOS Press, 2010.
- [29] M. Chang, M. He, and X. Luo. Astoncat-plus: an efficient specialist for the tac market design tournament. In *Proceedings of the Twenty-Second international joint conference on Artificial Intelligence-Volume Volume One*, pages 146–151. AAAI Press, 2011.
- [30] Meng Chang, Minghua He, and Xudong Luo. AstonCAT-Plus: A smart e-market agent in the competitive environment of cat multiple markets. *Journal of Autonomous Agents and Multi-Agent Systems*, 2012. (Submitted).
- [31] C. Chiarella. The dynamics of speculative behaviour. *Annals of operations research*, 37(1):101–123, 1992.
- [32] S. Cincotti, L. Ponta, and M. Raberto. A multi-assets artificial stock market with zero-intelligence traders. In *WEHIA 2005(13-15 June 2005)*, 2005.
- [33] D. Cliff. Minimal-intelligence agents for bargaining behaviours in market-based environments. *Technical Report HP-97-91, Hewlett-Packard Research Laboratories, Bristol, England*, 1997.
- [34] D. Cliff. Evolution of market mechanism through a continuous space of auction-types. In *Evolutionary Computation, 2002. CEC’02. Proceedings of the 2002 Congress on*, volume 2, pages 2029–2034. IEEE, 2002.

- [35] D. Cliff and J. Bruten. Minimal-intelligence agents for bargaining behaviours in market-based environments. *Technical Report HPL-97-91, Hewlett Packard Labs*, 1997.
- [36] K.J. Cohen, S.F. Maier, R.A. Schwartz, and D.K. Whitcomb. A simulation model of stock exchange trading. *Simulation*, 41(5):181–191, 1983.
- [37] N. Craswell, O. Zoeter, M. Taylor, and B. Ramsey. An experimental comparison of click position-bias models. In *Proceedings of the international conference on Web search and web data mining*, pages 87–94. ACM, 2008.
- [38] G. Creamer. Model calibration and automated trading agent for euro futures. *Quantitative Finance*, 12(4):531–545, 2012.
- [39] G. Creamer and Y. Freund. Automated trading with boosting and expert weighting. *Quantitative Finance*, 10(4):401–420, 2010.
- [40] G.H. Creamer and Y. Freund. A boosting approach for automated trading. *Journal of Trading*, 2(3):84–96, 2007.
- [41] M.G. Daniels, J.D. Farmer, L. Gillemot, G. Iori, and E. Smith. Quantitative model of price diffusion and market friction based on trading as a mechanistic random process. *Physical review letters*, 90(10):108–102, 2003.
- [42] A. Das, I. Giotis, A.R. Karlin, and C. Mathieu. On the effects of competing advertisements in keyword auctions. *Fourth Workshop on Ad Auction July 8-9, Chicago*, 2008.
- [43] R. Das, J.E. Hanson, J.O. Kephart, and G. Tesauero. Agent-human interactions in the continuous double auction. In B. Nebel, editor, *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI2001)*, pages 1169–1176, Seattle, WA, USA, 2001.
- [44] S. Das. A learning market-maker in the glosen–milgrom model. *Quantitative Finance*, 5(2):169–180, 2005.
- [45] K. Davis. ‘zero intelligence’ trading closely mimics stock market. *New Scientist*, 2005.
- [46] P. De Grauwe, H. Dewachter, and M. Embrechts. Exchange rate theory: chaotic models of foreign exchange markets. *Status: published*, 1995.
- [47] M.A.H. Dempster and C.M. Jones. A real-time adaptive trading system using genetic programming. *Quantitative Finance*, 1:397–413, 2000.

- [48] M.A.H. Dempster and V. Leemans. An automated fx trading system using adaptive reinforcement learning. *Expert Systems with Applications*, 30(3):543–552, 2006.
- [49] M.A.H. Dempster, T.W. Payne, Y. Romahi, and G.W.P. Thompson. Computational learning techniques for intraday fx trading using popular technical indicators. *IEEE Transactions on neural networks*, 12(4):744–754, 2001.
- [50] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 7(1):242–259, 2007.
- [51] N. Ehrentreich. A corrected version of the santa fe institute artificial stock market model. *Complexity*, 2003.
- [52] A. Eisingerich and T. Kretschmer. In e-commerce, more is more. *Harvard Business Review*, 86:20–21, 2008.
- [53] I. Erev and A.E. Roth. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American economic review*, pages 848–881, 1998.
- [54] D. Fain and J. Pedersen. Sponsored search: a brief history. *Bulletin of the American Society for Information Science and Technology*, 32(2):12–13, 2006.
- [55] J. Feng, H.K. Bhargava, and D.M. Pennock. Implementing sponsored search in web search engines: Computational evaluation of alternative mechanisms. *INFORMS Journal on Computing*, 19(1):137–148, 2007.
- [56] Y. Feng, R. Yu, and P. Stone. Two stock-trading agents: Market making and technical analysis. *Agent-Mediated Electronic Commerce V. Designing Mechanisms and Systems*, pages 18–36, 2004.
- [57] F. Franci, R. Marschinski, and L. Matassini. Learning the optimal trading strategy. *Physica A: Statistical Mechanics and its Applications*, 294(1):213–225, 2001.
- [58] J.A. Frankel and K.A. Froot. Explaining the demand for dollars: International rates of return and the expectations of chartists and fundamentalists. In R. Chambers and P. Paarlberg, editors, *Agriculture, Macroeconomics, and the Exchange Rate*. Westview Press, Boulder, CO., 1988.
- [59] D. Friedman. *The double auction market: institutions, theories, and evidence*, volume 14. Westview Press, 1993.
- [60] D. Friedman and C. Rich. The matching market institution. *Handbook of experimental economics results*, 1:109–114, 2008.

- [61] A. Ghose and S. Yang. An empirical analysis of search engine advertising: Sponsored search in electronic markets. *Management Science*, 55(10):1605–1622, 2009.
- [62] S. Gjerstad and J. Dickhaut. Price formation in double auctions. *Games and economic behavior*, 22(1):1–29, 1998.
- [63] D. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137, 1993.
- [64] A. Goldfarb and C. Tucker. Search engine advertising: Channel substitution when pricing ads to context. *Management Science*, 57(3):458–470, 2011.
- [65] B.A. Good. The changing face of money: will electronic money be adopted in the united states? *Taylor & Francis*, pages 80–81, 2000.
- [66] M. Grinblatt and S. Titman. Performance measurement without benchmarks: An examination of mutual fund returns. *Journal of Business*, 66(1):47–68, January 1993.
- [67] M. Grinblatt, S. Titman, and R. Wermers. Momentum investment strategies, portfolio performance, and herding: A study of mutual fund behavior. *The American Economic Review*, 85(5):1088–1105, December 1995.
- [68] S.J. Grossman and J.E. Stiglitz. On the impossibility of informationally efficient markets. *The American Economic Review*, pages 393–408, 1980.
- [69] F.A. Haight. Handbook of the poisson distribution. 1967.
- [70] L. Harris, G. Sofianos, and J. E. Shapiro. Program trading and intraday volatility. *The Review of Financial Studies*, 10(4):1035–1064, 1997.
- [71] M. He, N.R. Jennings, and H.F. Leung. On agent-mediated electronic commerce. *IEEE Transactions on Knowledge and Data Engineering*, 15(4):985–1003, 2003.
- [72] M. He, H. Leung, and N.R. Jennings. A fuzzy-logic based bidding strategy for autonomous agents in continuous double auctions. *IEEE Transactions on Knowledge and Data Engineering*, 15(6):1345–1363, 2003.
- [73] D.L. Hoffman and T.P. Novak. How to acquire customers on the web. *Harvard business review*, 78(3):179–188, 2000.
- [74] C.H. Hommes. Financial markets as non-linear adaptive evolutionary systems. *Quantitative Finance*, 1:149–167, 2001.
- [75] S. Honari, M. Ebadi, A. Foshati, M. Gomrokchi, J. Benatahr, and B. Khosravifar. Price estimation of persiancat market equilibrium. In *Workshop on Trading Agent Design and Analysis (TADA 2009), IJCAI*. Citeseer, 2009.

- [76] P.H. Hsu and C.M. Kuan. Re-examining the profitability of technical analysis with data snooping checks. *Journal of Financial Econometrics*, 3(4):606–628, 2005.
- [77] B.J. Jansen and T. Mullen. Sponsored search: An overview of the concept, history, and technology. *International Journal of Electronic Business*, 6(2):114–131, 2008.
- [78] B.J. Jansen and A Spink. The effect on click-through of combining sponsored and non-sponsored search engine results in a single listing. *Proceedings of 2007 Workshop on Sponsored Search Auctions*, 2007.
- [79] S. Jap. An exploratory study of the introduction of online reverse auctions. *The Journal of Marketing*, 2003.
- [80] N. Jegadeesh and S. Titman. Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of Finance*, 48(1):65–91, 2012.
- [81] N.R. Jennings. An agent-based approach for building complex software systems. *Communications of ACM*, 44(4):35–41, 2001.
- [82] N.R. Jennings, K. Sycara, and M. Wooldridge. A roadmap of agent research and development. *Autonomous agents and multi-agent systems*, 1(1):7–38, 1998.
- [83] T. Joachims, L. Granka, B. Pan, H. Hembrooke, F. Radlinski, and G. Gay. Evaluating the accuracy of implicit feedback from clicks and query reformulations in web search. *ACM Transactions on Information systems*, 25(2), 2007.
- [84] P. Jordan, B. Cassell, L. Callender, and M. Wellman. The ad auctions game for the 2009 trading agent competition. Technical report, University of Michigan, USA, 2009.
- [85] K.A. Kavajecz and E.R. Odders-White. Technical analysis and liquidity provision. *Review of Financial Studies*, 17(4):1043–1071, 2004.
- [86] M. Kearns and L. Ortiz. The penn-lehman automated trading project. *Intelligent Systems, IEEE*, 18(6):22–31, 2003.
- [87] D. Kempe and M. Mahdian. A cascade model for externalities in sponsored search. *Internet and Network Economics*, pages 585–596, 2008.
- [88] G. Kendall and Y. Su. The co-evolution of trading strategies in a multi-agent based simulated stock market through the integration of individual learning and social learning. *Proceedings of IEEE*, pages 2298–2305, 2003.
- [89] G. Kendall and Y. Su. A particle swarm optimisation approach in the construction of optimal risky portfolios. In *Proceedings of the 23rd IASTED International Multi-Conference on Applied Informatics*, pages 140–145. Citeseer, 2005.

- [90] D. Kilpatrick. Keyword advertising on a cost-per-click model, 2003.
- [91] G. Kim and H.M. Markowitz. Investment rules, margin, and market volatility. *The Journal of Portfolio Management*, 16(1):45–52, 1989.
- [92] K. Kim. Artificial neural networks with feature transformation based on domain knowledge for the prediction of stock index futures. *International Journal of Intelligent Systematic Accounting Financial Management*, 12(3), 2004.
- [93] A.P. Kirman. Epidemics of opinion and speculative bubbles in financial markets. In M. Taylor, editor, *Money and Financial Markets*. Macmillan, London, 1991.
- [94] A.P. Kirman and N.J. Vriend. Evolving market structure: An ace model of price dispersion and loyalty. *Journal of Economic Dynamics and Control*, 25(3):459–502, 2001.
- [95] B. Kitts and B. Leblanc. Optimal bidding on keyword auctions. *Electronic Markets*, 14(3):186–201, 2004.
- [96] D. Ladley. Zero intelligence in economics and finance. *The Knowledge Engineering Review*, pages 1–24, 2004.
- [97] S. Lahaie and D. Pennock. Revenue analysis of a family of ranking rules for keyword auctions. In *Proceedings of the Eighth ACM Conference on Electronic Commerce*, pages 50–56, San Diego, 2007.
- [98] S. Lahaie, D.M. Pennock, A. Saberi, and R.V. Vohra. Sponsored search auctions. *Algorithmic Game Theory*, pages 699–716, 2007.
- [99] B. LeBaron. Agent-based computational finance: Suggested readings and early research. *Journal of Economic Dynamics and Control*, 24(5):679–702, 2000.
- [100] B. LeBaron. Building the santa fe artificial stock market. *Physica A*, 2002.
- [101] B. LeBaron, W.B. Arthur, and R. Palmer. Time series properties of an artificial stock market. *Journal of Economic Dynamics and Control*, 23(9):1487–1516, 1999.
- [102] M. Lettau. Explaining the facts with adaptive agents: The case of mutual fund flows. *Journal of Economic Dynamics and Control*, 21(7):1117–1147, 1997.
- [103] H. Levy, M. Levy, and S. Solomon. *Microscopic simulation of financial markets: from investor behaviour to market phenomena*. Academic Press, 2000.
- [104] M. Levy, H. Levy, and S. Solomon. A microscopic model of the stock market: cycles, booms, and crashes. *Economics Letters*, 45(1):103–111, 1994.

- [105] M. Levy, S. Solomon, and RAM GIVAT. Dynamical explanation for the emergence of power law in a stock market model. *International Journal of Modern Physics C*, 7(01):65–72, 1996.
- [106] A.W. Lo, H. Mamaysky, and J. Wang. Foundations of technical analysis: Computational algorithms, statistical inference, and empirical implementation. *The Journal of Finance*, 55(4):1705–1770, 2002.
- [107] D. Lucking-Reiley. Vickrey auctions in practice: From nineteenth-century philately to twenty-first-century e-commerce. *The Journal of Economic Perspectives*, pages 183–192, 2000.
- [108] X. Luo, C. Miao, N.R. Jennings, M. He, Z. Shen, and M. Zhang. Kemnad: a knowledge engineering methodology for negotiating agent development. *Computational Intelligence*, 28(1):51–105, 2012.
- [109] Y. Luo, K. Liu, and D.N. Davis. A multi-agent decision support system for stock trading. *Network, IEEE*, 16(1):20–27, 2002.
- [110] T. Lux. Time variation of second moments from a noise trader/infection model. *Journal of Economic Dynamics and Control*, 22(1):1–38, 1997.
- [111] T. Lux. The socio-economic dynamics of speculative markets: interacting agents, chaos, and the fat tails of return distributions. *Journal of Economic Behavior & Organization*, 33(2):143–165, 1998.
- [112] T. Lux and M. Marchesi. Scaling and criticality in a stochastic multi-agent model of a financial market. *Nature*, 397(6719):498–500, 1999.
- [113] T. Lux and M. Marchesi. Volatility clustering in financial markets: A microsimulation of interacting agents. *International Journal of Theoretical and Applied Finance*, 3(04):675–702, 2000.
- [114] J.K. MacKie-Mason and M.P. Wellman. Automated markets and trading agents. *Handbook of computational economics*, 2:1381–1431, 2006.
- [115] A. Madhavan, M. Richardson, and M. Roomans. Why do security prices change? a transaction-level analysis of nyse stocks. *Review of Financial Studies*, 10(4):1035–1064, 1997.
- [116] S. Markose, E. Tsang, and S.M. Jaramillo. The red queen principle and the emergence of efficient financial markets: an agent based approach. *Nonlinear Dynamics and Heterogeneous Interacting Agents*, pages 287–303, 2005.

- [117] S. Martinez-Jaramillo. *Artifical Financial Markets: An Agent Based Approach to Reproduce Stylized Facts and to Study the Red Queen Effect*. PhD thesis, University of Essex, 2007.
- [118] R.P. McAfee and J. McMillan. Auction and bidding. *Journal of Economic Literature*, 25(2):699–738, 1987.
- [119] K. Moriyama, M. Matsumoto, K. Fukui, S. Kurihara, and M. Numao. Reinforcement learning on a futures market simulator. *Journal of Universal Computer Science*, 14(7):1136–1153, 2008.
- [120] M.R. Morris. A survey of collaborative web search practices. In *Proceedings of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, pages 1657–1660. ACM, 2008.
- [121] L. Muchnik and S. Solomon. Markov nets and the netlab platform: Application to continuous double auction. *Complexity Hints for Economic Policy*, pages 157–180, 2007.
- [122] S. Muthukrishnan, M. Pal, and Z. Svitkina. Stochastic models for budget optimization in search-based advertising. *Internet and Network Economics*, pages 131–142, 2007.
- [123] M. Naldi, A. Pavignani, A. Grillo, A. Lentini, and G.F. Italiano. The competitor busting strategy in keyword auctions: Who’s worst hit? In *Communication Networks and Services Research Conference (CNSR), 2011 Ninth Annual*, pages 153–160. IEEE, 2011.
- [124] C. Neely, P. Weller, and R. Dittmar. Is technical analysis in the foreign exchange market profitable? a genetic programming approach. *Journal of Financial and Quantitative Analysis*, 32(4), 1997.
- [125] C.J. Neely. Technical analysis in the foreign exchange market: a layman’s guide. *Federal Reserve Bank of St. Louis Review*, pages 23–38, 1997.
- [126] S.N. Neftci. Naive trading rules in financial markets and wiener-kolmogorov prediction theory: A study of “technical analysis”. *Journal of Business*, pages 549–571, 1991.
- [127] Y. Nevmyvaka, Y. Feng, and M. Kearns. Reinforcement learning for optimized trade execution. In *Proceedings of the 23rd international conference on Machine learning*, pages 673–680. ACM, 2006.
- [128] J. Nicolaisen, V. Petrov, and L. Tesfatsion. Market power and efficiency in a computational electricity market with discriminatory double-auction pricing. *Evolutionary Computation, IEEE Transactions on*, 5(5):504–523, 2001.

- [129] J. Niu, K. Cai, P. McBurney, and S. Parsons. An analysis of entries in the first tac market design competition. In *Web Intelligence and Intelligent Agent Technology, 2008. WI-IAT'08. IEEE/WIC/ACM International Conference on*, volume 2, pages 431–437. IEEE, 2008.
- [130] J. Niu, K. Cai, S. Parsons, E. Gerding, and P. McBurney. Characterizing effective auction mechanisms: Insights from the 2007 tac market design competition. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 2*, pages 1079–1086. International Foundation for Autonomous Agents and Multiagent Systems, 2008.
- [131] 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*, pages 1649–1650. International Foundation for Autonomous Agents and Multiagent Systems, 2008.
- [132] J. Niu, K. Cai, S. Parsons, P. McBurney, and E.H. Gerding. What the 2007 TAC market design game tells us about effective auction mechanisms. *Autonomous Agents and Multi-Agent Systems*, 21(2):172–203, 2010.
- [133] J. Niu, K. Cai, S. Parsons, and E. Sklar. Reducing price fluctuation in continuous double auctions through pricing policy and shout improvement. In *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, pages 1143–1150. ACM, 2006.
- [134] J. Niu, K. Cai, S. Parsons, and E. Sklar. Some preliminary results on competition between markets for automated traders. In *Proceedings of AAAI-07 Workshop on Trading Agent Design and Analysis (TADA-07), Vancouver, Canada, 2007*.
- [135] J. Niu, A. Mmoloke, P. McBurney, and S. Parsons. CATP a communication protocol for CAT games version 2.0. Technical report, Technical Report ULCS-09-012, Department of Computer Science, University of Liverpool, Liverpool, UK, 2009.
- [136] M. O'Neill. Did beezid pay too much for their lindsay lohan house arrest promo?, 2011.
- [137] M.J. Osborne and A. Rubinstein. *A course in game theory*. MIT press, 1994.
- [138] A. O'sullivan and S.M. Sheffrin. *Economics: Principles in action*. Boston, Mass.: Pearson/Prentice Hall, 2007.
- [139] D. Pardoe and P. Stone. Developing adaptive auction mechanisms. *ACM SIGecom Exchanges*, 5(3):1–10, 2005.

- [140] S. Parsons, M. Marcinkiewicz, J. Niu, and S. Phelps. Everything you wanted to know about double auctions, but were afraid to (bid or) ask. *Technical Report, Department of Computer & Information Science, Brooklyn College, City University of New York*, 2006.
- [141] S. Parsons, J.A. Rodriguez-Aguilar, and M. Klein. Auctions and bidding: A guide for computer scientists. *ACM Computing Surveys*, 43(2):10, 2011.
- [142] A. Pekeč and M.H. Rothkopf. Combinatorial auction design. *Management Science*, 49:1485–1503, November 2003.
- [143] J.M. Perloff. *Microeconomics*. Addison Wesley, 1998.
- [144] S. Phelps. JASA - Java Auction Simulation API. <http://jasa.sourceforge.net/>, 2005.
- [145] S. Phelps, S. Parsons, E. Sklar, and P. McBurney. Using genetic programming to optimise pricing rules for a double auction market. In *Proceedings of the workshop on Agents for Electronic Commerce*. Citeseer, 2003.
- [146] K. Qureshi, B. Nazir, and MA Shah. Markup based continuous double auction for resource allocation in market grid. *Engineering e-Transaction (ISSN 1823-6379)*, 6(1):50–54, 2011.
- [147] M. Raberto, S. Cincotti, S.M. Focardi, and M. Marchesi. Agent-based simulation of a financial market. *Physica A: Statistical Mechanics and its Applications*, 299(1):319–327, 2001.
- [148] A. Rapoport, H. Otsubo, B. Kim, and W.E. Stein. Unique bid auctions: Equilibrium solutions and experimental evidence. *MPRA Paper*, 4185, 2007.
- [149] E. Robinson, P. McBurney, and X. Yao. How specialised are specialists? generalisation properties of entries from the 2008 tac market design competition. In *Proceedings of the International Trading Agent Design and Analysis Workshop (TADA 2009), Pasadena, CA, USA*, 2009.
- [150] B.R. Routledge. Genetic algorithm learning to choose and use information. *Macroeconomic dynamics*, 5(2):303–325, 2001.
- [151] P. Rusmevichientong and D.P. Williamson. An adaptive algorithm for selecting profitable keywords for search-based advertising services. In *Proceedings of the 7th ACM conference on Electronic commerce*, pages 260–269. ACM, 2006.
- [152] O. Rutz and R.E. Bucklin. A model of individual keyword performance in paid search advertising. *Available at SSRN 1024765*, 2007.

- [153] W.H. Sandholm. Negative externalities and evolutionary implementation. *Review of Economic Studies*, 72(3):885–915, 2005.
- [154] W.H. Sandholm. Pigouvian pricing and stochastic evolutionary implementation. *Journal of Economic Theory*, 132(1):367–382, 2007.
- [155] A. Shah. Competing exchanges: The international dimension. *Business Standard*, 1997.
- [156] D. Sherrington and T. Galla. The minority game: effects of strategy correlations and timing of adaptation. *Physica A: Statistical Mechanics and its Applications*, 324(1):25–29, 2003.
- [157] J. Sohn, S. Lee, and T. Mullen. Attracting intra-marginal traders across multiple markets. In *Workshop on Trading Agent Design and Analysis (TADA 2009)*, 2009.
- [158] L.C. Stavrogiannis and P.A. Mitkas. CAT 2008 post-tournament evaluation: The mertacor’s perspective. *Association for the Advancement of Artificial Intelligence*, 2009.
- [159] P. Stone and A. Greenwald. The first international trading agent competition: Autonomous bidding agents. *Electronic Commerce Research*, 5(2):229–265, 2005.
- [160] S. Sunder. Market as an artifact aggregate efficiency from zero intelligence traders. In M.E. Augier and J.G. March, editors, *Models of a Man: Essays in Memory of Herbert*, pages 501–519, Cambridge, MA, 2004. MIT Press.
- [161] R.S. Sutton and A.G. Barto. *Reinforcement learning: An introduction*. Cambridge University Press, 1998.
- [162] F.E.H. Tay and L. Cao. Application of support vector machines in financial time series forecasting. *Omega*, 29(4):309–317, 2001.
- [163] M.P. Taylor and H. Allen. The use of technical analysis in the foreign exchange market. *Journal of international Money and Finance*, 11(3):304–314, 1992.
- [164] R. Telang, P. Boatwright, and T. Mukhopadhyay. A mixture model for internet search-engine visits. *Journal of Marketing Research*, pages 206–214, 2004.
- [165] G. Tesauro and J.L. Bredin. Strategic sequential bidding in auctions using dynamic programming. In *Proceedings of the first international joint conference on Autonomous agents and multiagent systems: part 2*, pages 591–598. ACM, 2002.
- [166] G. Tesauro and R. Das. High-performance bidding agents for the continuous double auction. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, pages 206–209. ACM, 2001.

- [167] L. Tesfatsion. Introduction to the special issue on agent-based computational economics. *Journal of Economic Dynamics and Control*, 25(3):281–293, 2001.
- [168] L. Tesfatsion. Agent-based computational economics: Growing economies from the bottom up. *Artificial life*, 8(1):55–82, 2002.
- [169] R.R. Trippi and D. DeSieno. Trading equity index futures with a neural network. *The Journal of Portfolio Management*, 19(1):27–33, 1992.
- [170] E.P.K. Tsang and S. Martinez-Jaramillo. Computational finance. *IEEE Computational Intelligence Society Newsletter*, 3(8), 2004.
- [171] H.R. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163–1178, 2007.
- [172] I. Veryzhenko, P. Mathieu, and O. Brandouy. Key points for realistic agent-based financial market simulations. In *Proceedings of ICAART (3rd Conf.)*, volume 2, pages 74–83, 2011.
- [173] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.
- [174] Y. Vorobeychik. Simulation-based game theoretic analysis of keyword auctions with low-dimensional bidding strategies. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pages 583–590. AUAI Press, 2009.
- [175] Y. Vorobeychik and D.M. Reeves. Equilibrium analysis of dynamic bidding in sponsored search auctions. *International Journal of Electronic Business*, 6(2):172–193, 2008.
- [176] P. Vytelingum. *The structure and behaviour of the Continuous Double Auction*. PhD thesis, University of Southampton, 2006.
- [177] P. Vytelingum, I.A. Vetsikas, B. Shi, and N.R. Jennings. IAMwildCAT: The Winning Strategy for the TAC Market Design Competition. In *Proceedings of 18th European Conference on Artificial Intelligence*, pages 428–434, 2008.
- [178] V. Walia, A. Byde, and D. Cliff. Evolving market design in zero-intelligence trader markets. In *E-Commerce, 2003. CEC 2003. IEEE International Conference on*, pages 157–164. IEEE, 2003.
- [179] W.E. Walsh, R. Das, G. Tesauero, and J.O. Kephart. Analysing complex strategic interactions in multi-agent games. In P. Gmytrasiewicz and S. Parsons, editors, *Proceedings of the AAAI Workshop on Game-Theoretic and Decision-Theoretic Agents (GTDT 2002)*, pages 206–209, Edmonton, Canada, 2002.

- [180] M.P. Wellman. A market-oriented programming environment and its application to distributed multicommodity flow problems. *arXiv preprint cs/9308102*, 1993.
- [181] M.P. Wellman, S.F. Cheng, D.M. Reeves, and K.M. Lochner. Trading agents competing: Performance, progress, and market effectiveness. *Intelligent Systems, IEEE*, 18(6):48–53, 2003.
- [182] D. Witkowska and E. Marcinkiewicz. Construction and evaluation of trading systems: Warsaw index futures. *International Advances in Economic Research*, 11(1):83–92, 2005.
- [183] M. Wooldridge. *An introduction to multiagent systems*. Wiley, 2002.
- [184] M. Wooldridge and N.R. Jennings. Intelligent agents: Theory and practice. *Knowledge Engineering Review*, 10(2):115–152, 1995.
- [185] P.R. Wurman. Dynamic pricing in the virtual marketplace. *IEEE Internet Computing*, 5:36–42, March/April 2001.
- [186] P.R. Wurman, W.E. Walsh, and M.P. Wellman. Flexible double auctions for electronic commerce: Theory and implementation. *Decision Support Systems*, 24(1):17–27, 1998.
- [187] P.R. Wurman, M.P. Wellman, and W.E. Walsh. A parametrization of the auction design space. *Games and economic behavior*, 35(1):304–338, 2001.
- [188] J. Yang. The efficiency of an artificial double auction stock market with neural learning agents. *Studies in Fuzziness and Soft Computing*, 100:85–106, 2002.
- [189] T. Yu and S.H. Chen. Using genetic programming with lambda abstraction to find technical trading rules. Technical report, Society for Computational Economics, 2004.
- [190] D. Zhao, D. Zhang, and L. Perrussel. Mechanism design for double auctions with temporal constraints. In *Proceedings of the Twenty-Second international joint conference on Artificial Intelligence-Volume Volume One*, pages 472–477. AAAI Press, 2011.
- [191] Y. Zhou, D. Chakrabarty, and R. Lukose. Budget constrained bidding in keyword auctions and online knapsack problems. *Internet and Network Economics*, pages 566–576, 2008.
- [192] Y. Zhou and R. Lukose. Vindictive bidding in keyword auctions. In *Proceedings of the ninth international conference on Electronic commerce*, pages 141–146. ACM, 2007.

Appendix A

Parameter Settings of AstonTAC and TAC AA competition

Table A.1: Parameter Settings of AstonTAC and TAC AA competition. Here we only list important parameters to the design of agent AstonTAC. For comprehensive information about TAC AA game parameter settings, refer to Table 1 in [84].

Symbol	Denotation	Setting
Game standard parameters		
D	Length of game	60 days
N	Number of advertising agents	8 agents
N^{High}	High capacity agents in a game	2 agents
N^{Med}	Medium capacity agents in a game	4 agents
N^{Low}	Low capacity agents in a game	2 agents
M	Number of simulated search users	90,000 users
K	Ad slot	5 slots
W	Aggregation window size for distribution capacity	5
USP	Unit sale profit	\$ 10
TE	Target effect	0.5
CSB	Component specialist bonus	0.5
MSB	Manufacturer specialist bonus	0.5
Pr_{burst}	Probability of searching burst	0.1
C^{cap}	Distribution capacity	300,400,500
λ	Distribution capacity discount	0.995
π_l	Baseline conversion rate	0.1,0.2,0.3
$P_{conversion,def}$	Baseline conversion rate	0.1, 0.2, 0.3
χ	Squashing parameter	$0 \leq \chi^4 \leq 1$
Parameters defined for AstonTAC		
$h_{v,q}$	Query value adapter of query q	$80\% \leq h_v \leq 120\%$
h_c	Distribution capacity adapter for day 0 and 1	$h_c \in \{85\%, 100\%, 115\%\}$
$r_{discount}$	Discount ratio for static MVPC	0.775

Appendix B

List of Symbols Used in Chapter 3

Table B.1: List of symbols used in Chapter 3.

Symbol	Denotation
CPC	Cost Per Click
CTR	Click Through Rate
PPC	Profit Per Conversion
MVPC	Market-based Value Per Click
$v_{static,q}$	Static MVPC of query q
$v_{dynamic,q}$	Dynamic MVPC of query q
q_c	Component part of query q
q_m	Manufacturer part of query q
$P_{conversion}$	Conversion rate in general
$P_{conversion,q}$	Conversion rate of query q
$P_{conversion,t}$	Timely conversion rate
$P_{conversion,crit}$	Critical timely conversion rate
$P_{conversion,std}$	Average baseline conversion rate
$P_{conversion,q,d+1}$	Estimated conversion rate of query q on day $d + 1$
$v_{con,q}$	Conversion value of query q
δ	Distribution capacity adapter
β_q	Ranking mechanism adapter for query q
$b_{0,q}$	Bid for query q in Phase One
b_q	Bid for query q in Phase Two
C_{crit}	Critical number of conversion
$C_{w,d+1}$	Conversion allowance for day $d + 1$ where w is the aggregate window
$c_{agg,d}$	W days aggregate conversion by day d
c_{d-4}	Number of conversions of all queries on day $d - 4$
$c_{q,d+1}$	Estimated conversion from query q on day $d + 1$
e_q	Estimated click through probability by the publisher for query q
e'_q	Estimated e_q
t_q	Ad display type of query q
$P_{click,q,d+1}$	Click probability on day $d + 1$ for query q
v'_{con}	General conversion value with respect to all queries
c'_{click}	General cost per click with respect to all queries
$impression_{q,i}$	Number of impression from query q day i
$revenue_{q,i}$	Revenue from query q on day i
$conversion_{q,i}$	Number of conversion from query q on day i
$click_{q,i}$	Number of clicks from query q on day i
$cost_{q,i}$	Cost from query q on day i

Appendix C

Parameter Settings and Symbol Denotation of AstonCAT-Plus

Table C.1: Parameter settings and symbol denotation of AstonCAT-Plus - Part 1.

Symbol	Denotation	Setting
p_s	Short-term equilibrium price	Dynamic
p_l	Long-term equilibrium price	Dynamic
\hat{p}^*	Estimated local market equilibrium price	Dynamic
W_{short}	Short window size for short-term equilibrium price	5
W_{long}	Long window size for long-term equilibrium price	20
ω_s	Weight of p_s in calculation of \hat{p}^*	Dynamic
\underline{b}	Minimum transacted bid	Dynamic
\bar{a}	Maximum transacted ask	Dynamic
α	Slack rate for short accepting threshold	Dynamic
$\alpha_{0,s}$	Initial α for ask	Dynamic
$\alpha_{0,b}$	Initial α for bid	Dynamic
β	Parameter used to flatten result of initial α	4
τ_s	Ask accepting threshold	Dynamic
τ_b	Bid accepting threshold	Dynamic
l_s	Ask accepting threshold ratio limit	1.05
l_b	Bid accepting threshold ratio limit	0.95
ρ	Transaction profit per transaction (TPT)	Dynamic
$\tilde{\rho}$	Average TPT	Dynamic
θ_l	Minimum TPT threshold for extra-marginal matches	$0.16(\bar{b}_t - \underline{a}_t)$
θ_s	Minimum TPT threshold for intra-marginal matches	$0.02(\bar{b}_t - \underline{a}_t)$
\bar{b}_t	Highest attempted bid	Dynamic
\underline{a}_t	Lowest attempted ask	Dynamic

Table C.2: Parameter settings and symbol denotation of AstonCAT-Plus - Part 2.

Symbol	Denotation	Setting
n_{match}	Number of matched bid-ask pairs	Dynamic
n_{trader}	Number of traders	Game setting
n_{market}	Number of market specialists	Game setting
n_{tar}	Trader target	Dynamic
n_{cur}	Average trader quantity with AstonCAT-Plus in last 15 days	Dynamic
$\bar{n}_{traders}$	All time mean of AstonCAT-Plus' daily trader	Dynamic
r_t	Market trend ratio benchmark	Dynamic
N/A	Up market trend threshold	1.16
N/A	Down market trend threshold	0.92
r_v	Moving trend ratio benchmark	Dynamic
N/A	Up moving trend threshold	1.006
N/A	Down moving trend threshold	0.97
N/A	Mini fee updating step	0.00025
N/A	Small fee updating step	0.0005
N/A	Large fee updating step	0.001
r_t	Market trend ration benchmark	Dynamic
$\tilde{\rho}_t$	Average trader profit	Dynamic
$\tilde{\rho}_r$	Average trader profit per transaction	Dynamic
\tilde{n}_t	Average transaction quantity per trader	Dynamic
φ	Allocative efficiency	Dynamic
ψ	Convergence coefficient	Dynamic
λ	Side-balance rate of traders	Dynamic

Appendix D

Table of Experiment Outcomes

Table D.1: CTR and profit of TacTex and AstonTAC in controlled experiment C of Chapter 3.

Capacity	Iteration	TacTex CTR	AstonTAC CTR	TacTex Profit	AstonTAC Profit
G300	608	10.37%	10.33%	40259.59	39760.22
	610	18.74%	13.13%	41214.53	47060.76
	611	9.71%	7.53%	41056.16	48381.48
	612	17.52%	9.65%	49091.12	41341.47
	613	8.73%	8.90%	33669.33	39504.69
G400	632	14.65%	10.61%	56906.73	43360.63
	633	5.23%	8.79%	46609.97	51228.60
	634	6.18%	11.13%	36640.32	54055.40
	635	9.37%	9.99%	55364.70	44375.77
	636	8.04%	9.69%	49942.48	47273.09
G500	637	6.79%	17.04%	53846.75	64848.68
	638	8.09%	10.74%	53988.70	57987.52
	639	9.89%	10.04%	47826.71	48354.27
	640	8.24%	16.97%	50209.70	65832.45
	641	8.90%	13.33%	52413.63	62823.46