

UNIVERSITY OF WESTERN SYDNEY
&
UNIVERSITY OF TOULOUSE

Mechanism Design for Dynamic Double Auctions

by
Dengji Zhao

A thesis submitted in partial fulfillment for the
degree of Doctor of Philosophy

in the
Intelligent Systems Lab
School of Computing, Engineering and Mathematics
&
Institut de Recherche en Informatique de Toulouse (IRIT)

27 June 2012

©2012 - Dengji Zhao

All rights reserved.

Declaration of Authorship

I, Dengji Zhao, declare that this thesis titled, ‘Mechanism Design for Dynamic Double Auctions’ and the work presented in it are my own. I confirm that:

- This work was done wholly while in candidature for a research degree at the two universities.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at the two universities or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Signed:

Date:

Mechanism Design for Dynamic Double Auctions

Abstract

A mechanism is a specification for the determination of economic decisions based on the information that is known by the individuals within the economy. Mechanism design is the discipline of designing mechanisms that lead to socially desirable outcomes in a context in which individuals are self-interested. Traditionally, mechanism design has focused on static settings in which the individuals (participants) are known to the mechanism prior to any decision being made. However, many real environments are dynamic, such as in a stock exchange where participants are arriving and departing at different times, and existing solutions for static settings are inappropriate.

Although mechanism design for dynamic settings has gained the attention of many researchers over the past decade, most of them have focused on one-sided dynamic market models; that is, either the supply or the demand of the market is dynamic, but not both. Online double auctions, in which the dynamics are two-sided, represent the dominant type of exchange market, but only limited studies have been conducted for online double auctions, due to the complexity of the dynamics. In order to address this gap, this thesis attacks the design problem in two types of online double auction: one type is decision-independent, where each trader's private information (that is, type) is observed independently and therefore cannot be changed by the decisions of the auction, and the other is decision-dependent, where each trader's private information depends on other traders and also varies in response to the decisions of the auction.

For the first type, this thesis studies a model in which each trader participates (or is active) in the market for only one period of time and the trader's valuation does not change during this period. First, it provides a computationally efficient optimal (offline) solution that is truthful, efficient and individually rational. This optimal solution is one kind of Vickrey-Clarke-Groves (VCG) mechanism, but it is computationally more efficient than the classical VCG mechanism. Apart from

servicing the online double auction design, this VCG mechanism also provides a dedicated solution for real trading environments such as futures exchanges. Next, it proposes a reduction framework within which to build an online/dynamic double auction by reducing it to an online one-sided auction. This reduction framework is notable because (1) well-studied online one-sided auctions can be easily reused, and (2) the key properties of the reduced online double auction match those of the online one-sided auction, which is very difficult even for static double auctions. In addition, this thesis shows that in this model it is impossible to design a deterministic online double auction that is truthful, individually rational and competitive for efficiency, although it shows that this is possible under certain assumptions.

The second type is approached by means of two steps. First, the double auction design problem is studied in an environment in which traders' valuations vary with respect to the number of units they trade, but without consideration for the dynamic nature of arrivals and departures. This environment can be mapped to the tremendously growing online shopping model, which leverages group buying, and in this thesis it is modelled as a multi-unit double auction. The thesis provides new insights (impossibilities and possibilities) into the design of multi-unit double auctions under group buying. In particular, it demonstrates that there are no budget-balanced, individually rational and truthful allocations that can guarantee a reasonable transaction size. In the second step, a more complex dynamic environment is envisaged, in which traders dynamically arrive and depart and their valuations change over time. This environment can be mapped to real stock exchanges. Since the models of traders in this kind of environment have been well studied in economics, this thesis addresses the auction design problem directly, based on these well-studied models. However, because traders' types are dependent on each other as well as on the decisions of the auction, a good auction also needs to learn traders' behaviours in order to make appropriate decisions in different environments. To that end, an auction design framework based on trader behaviours is developed. This framework demonstrates how auctions can be designed to analyse market dynamics (or trader behaviours) and then use trader behaviours to guide market decisions such that desirable resource allocations are achievable by, for example, attracting more good traders to return to the market in the future.

Acknowledgements

Before I went to university I never dreamed that I would have what I have now. To reach this point, I owe much to a great many people. I deeply appreciate their kindness, support and guidance.

I would like to express my gratitude to the following:

To my co-supervisors, Dongmo Zhang and Laurent Perrussel. Apart from the many things you have taught me about research, your inspiration, encouragement and willingness to help provided me with a great learning environment. To the rest of my supervisory panel, Yan Zhang and Andreas Herzig, I very much appreciate your care as I worked to complete my thesis.

To my Mum, who is my hero and who taught me a much about life. I cannot imagine how difficult your life has been after losing my Dad and my sister when I was 10 years old. Your attitude towards life and your endless support to the family drove me towards finding a meaningful future. To my grandma, who has been my guiding light throughout my life. To my stepfather, thank you for joining the family; I cannot go far without you. To my extended family, especially my uncle Yushun Li and my cousin Yanfang Peng, for your love and support.

To FAN Xuhe, HUANG Gang, LUO Wenbin, WANG Lian, WONG Ka Weng, KWONG Ying Wa, CHEANG Chak Fong, TANG Zesheng, XU Aoao, Michael Thielscher, Stephan Schiffel, Frank Ciesinski and Monika Mayer, who have guided and influenced me throughout my life. I am extremely fortunate to have you.

To my dearest friends, Yao Xu and Yonghua Chen: I just could not make all this happen without you. Also to my dear friends David Black, Minyi Li, Christy Liang, Ji Ruan, Yunjing Xiang, Heng Zhang, Zhiqiang Zhuang, Ben Gu and his family, whose help, caring and sharing have brought me an enjoyable PhD journey.

To my friends from the two great research labs, ISL and IRIT, especially to Chun Gao and Md Khan from ISL for the wonderful times we had in the *jackaroo* team, to Yun Bai, Leanne Ryland, Weixing Zheng, Yi Zhou, Sylvie Doutre and Jean-Marc Thévenin for your support and encouragement. Also to the AI community, especially Enrico Gerding, Vincent Conitzer, Lirong Xia, Ruggiero Cavallo, David Pennock, Wolf Ketter, Yingqian Zhang, Andreas Symeonidis and Tim Miller for your support and stimulating suggestions.

Finally, to the two universities and the Australian Research Council (ARC) for the funding and support provided for my PhD research. This work was supported by the ARC Discovery project: Logical Foundation and Implementation Technology for Automated Negotiation (DP0988750). The Cotutelle joint supervision program by the two universities, with the help of my co-supervisors, has provided me with an excellent learning experience.

Contents

Declaration of Authorship	i
Abstract	ii
Acknowledgements	iv
List of Figures	ix
List of Tables	x
Abbreviations	xi
1 Introduction	1
1.1 Motivation	1
1.1.1 Many Dynamic Environments Require Better Solutions	1
1.1.2 Existing Solutions are NOT Sufficient	2
1.1.3 Online Double Auction Design is More Challenging	3
1.2 Methodologies	4
1.2.1 Proposed Approaches	4
1.2.2 Existing Approaches	5
1.3 Major Contributions	6
1.4 Related Work	8
1.4.1 Static Demand and Dynamic Supply	9
1.4.2 Dynamic Demand and Static Supply	9
1.4.3 Dynamic Demand and Dynamic Supply	10
1.5 Outline of Chapters	11
2 Double Auction Models: Static and Dynamic	13
2.1 Design Objectives	14
2.2 Static Double Auctions	15
2.2.1 The Model	16
2.3 Approaching Online Double Auctions	18

2.3.1	The Model	19
2.3.1.1	Double Auction with Temporal Constraints	19
2.3.1.2	Double Auction under Group Buying	20
2.4	Online Double Auctions	22
2.4.1	The Model	23
2.4.1.1	Decision-independent Dynamic Environment	23
2.4.1.2	Decision-dependent Dynamic Environment	24
2.5	Summary	24
3	Matching in a Double Auction	26
3.1	Static Matching	26
3.1.1	Equilibrium Matching	27
3.1.1.1	Impossibility	27
3.1.2	Maximal Matching	29
3.1.2.1	Complexity Analysis	30
3.2	Online Matching	31
3.2.1	The Ranking Algorithm (for One-sided Dynamics)	31
3.2.2	The Greedy Algorithm (for Two-sided Dynamics)	33
3.3	Summary	34
4	Double Auction with Temporal Constraints	35
4.1	Introduction	36
4.2	The Model	37
4.3	Graph Representation	40
4.4	Efficient and Truthful Policy Design	44
4.4.1	Efficient & Monotonic Allocation Policy	44
4.4.2	Truthful Payment Policy	47
4.4.3	Computational Complexity	52
4.5	Summary	53
5	Multi-unit Double Auction under Group Buying	55
5.1	Introduction	56
5.2	The Model	59
5.3	A BB, IR and Buyer-truthful MDA	63
5.4	A BB, IR and Seller-truthful MDA	65
5.5	Existence of (W)BB, IR and Truthful MDAs	68
5.5.1	Competitive MDAs	73
5.6	Summary	74
6	Online Double Auction	77
6.1	Introduction	78
6.2	Preliminaries and Notations	80
6.3	No Deterministic Online Double Auctions are Competitive	82
6.4	A Deterministic & Competitive Online Double Auction	84
6.4.1	Specification of \mathcal{M}_{greedy}	84

6.4.2	Properties of \mathcal{M}_{greedy}	87
6.4.2.1	Truthfulness	87
6.4.2.2	Efficiency	91
6.5	Reducing Online Double Auctions to Online One-sided Auctions	95
6.5.1	The Reduction	96
6.5.2	Key Properties of \mathcal{M}_A	98
6.6	Summary	101
7	Behaviour-based Adaptive Double Auction	103
7.1	Introduction	104
7.2	Preliminary	106
7.3	Behaviour-based Trader Classification	107
7.3.1	Data Acquisition	108
7.3.2	Defining Categories of Trader	108
7.3.3	Category Recognition from Trader's Behaviour	109
7.3.3.1	Trading Time Distribution	109
7.3.3.2	Stability vs Intra-marginality	110
7.4	Behaviour-based Policy Design	111
7.4.1	A Search Space of Behaviour-based Policies	111
7.4.1.1	Accepting Policy	111
7.4.1.2	Matching Policy	112
7.4.1.3	Pricing Policy	113
7.4.1.4	Clearing Policy	114
7.4.1.5	Charging Policy	115
7.4.2	Searching Adaptive Mechanisms	115
7.4.2.1	Trader-dependent Mechanism Design	116
7.4.2.2	Adaptive Mechanisms with Trader-dependent Mechanisms	116
7.4.3	Experiments	117
7.5	A Framework for Behaviour-based Mechanism Design	119
7.6	Summary	120
8	Conclusion and Future Work	121
8.1	Summary of the Major Contributions	121
8.2	Future Work	123
	A Published Work	125
	Bibliography	126

List of Figures

3.1	Equilibrium Matching vs Maximal Matching	30
4.1	Example of Alternating Paths	42
4.2	Reachability Example	51
4.3	MBM Allocation and MM Payment	52
6.1	Running Example of \mathcal{M}_{greedy}	85
6.2	Seller Manipulation Examples I	90
6.3	Seller Manipulation Examples II	90
6.4	Best-first Allocation vs Optimal Allocation	91
6.5	A Special Case of Best-first Allocation	92
6.6	A General Case of Best-first Allocation	95
6.7	A Running Example of $\mathcal{M}_{\mathcal{A}}$	98

List of Tables

7.1	GD Attractive Mechanism	118
7.2	Average Trading Time Distribution of Each Type of Trader	118
7.3	Average Trading Time Distribution of Buyers	119

Abbreviations

DA	D ouble A uction
ODA	O nline D ouble A uction
MDA	M ulti-unit D ouble A uction
IC	I ncentive C ompatibility
(W)BB	(W eakly) B udget B alance
IR	I ndividual R ationality

*To my mum, Dingjian Cui, my dad, Yexuan Zhao, my
stepfather Zhongwen Zhang and my grandma*

Chapter 1

Introduction

1.1 Motivation

1.1.1 Many Dynamic Environments Require Better Solutions

A mechanism is a specification for the determination of economic decisions based on the information that is known by the individuals within the economy [Myerson, 2008a]. Mechanism design is the discipline of designing mechanisms (or information games) that lead to socially desirable outcomes in a context in which individuals are self-interested, and hold private information (which is known as *type*). Participants of a mechanism are asked to report their private information to the mechanism which selects an outcome that satisfies some properties.

Traditionally, mechanism design has focused on static settings in which all private information required for future decisions is known to the mechanism (or the decision-maker) at the start, and normally decisions are made at a single point in time. For instance, the Vickrey auction is designed for static environments, where there is one item for sale, each buyer submits his or her willing payment

(or valuation) for the item and the buyer who is willing to pay the highest price is awarded the right to purchase the item at the price offered by the second-highest bidder [Vickrey, 1961]. The Vickrey auction requires that all buyers have to be available at a specific time so that the auctioneer is able to collect sufficient information for its decision-making, otherwise the desired properties of the auction will not hold. However, auctions in many environments are dynamic; for example, at the New York Stock Exchange participants are arriving and departing at varying times and the market owner (or the mechanism) has to make a sequence of decisions over time rather than at a single point in time.

1.1.2 Existing Solutions are NOT Sufficient

Since there are many environments are dynamic, mechanism design for dynamic environments is necessary. Additionally, existing solutions for static settings are insufficient in dynamic environments. For instance, a seller is selling a house, and each buyer comes at a different time with a willing payment to buy the house and a waiting period during which the seller has to decide whether or not to sell the house to this buyer. In this situation, the Vickrey auction does not work properly, because the seller does not know if the willing payment of a buyer is the highest until all buyers have arrived. Unfortunately, the seller also cannot wait until all buyers have arrived, as the buyer with highest willing payment might have already left at that time. Apart from the challenge posed by the uncertainty about participants, the decision-making of a mechanism in a dynamic environment is also challenged by participants' strategical play with their arrival and departure; for example, a participant available/arriving at time t_1 might not report to the market until $t_2 > t_1$ if it is in his or her interest to do so. Existing solutions for static settings cannot handle this kind of manipulation considering the dynamic nature of arrivals and departures. Moreover, the private information of participants in a

dynamic environment might change over time and therefore, locally best solutions from a static auction's point of view might not be optimal over time.

1.1.3 Online Double Auction Design is More Challenging

An environment is dynamic if its participants are arriving and departing over time [Parkes, 2007], the valuation of each participant is changing over time [Bergemann and Välimäki, 2010, Cavallo and Parkes, 2008], or both [Cavallo et al., 2009]. Mechanism design for dynamic environments has gained the attention of many researchers over the past decade, but most of them have focused on one-sided dynamic market models; that is, either the supply or the demand of the market is dynamic, but not both [Parkes, 2007]. The house selling example in the previous section is the case in which only the demand is dynamic. When we consider a situation in which both supply and demand are dynamic, we normally talk about *online double auction* where multiple sellers and multiple buyers trade a commodity at any time they wish. An online double auction has to match sellers and buyers dynamically and calculate a payment for each matched trader without the knowledge of traders/orders coming afterwards; that is, without the benefit of hindsight about future traders and/or types. Such uncertainty is more challenging for double auction design because the modelling of traders' bidding behaviour in double auctions is 'immensely complicated' even in a static case [McAfee, 1992]. This is because both buyers and sellers are playing strategically in double auctions and the auctioneer (or the market owner) has no control of either side. Due to the complexity of the dynamics, only limited studies have been conducted for online double auctions [Blum et al., 2006, Bredin et al., 2007]. However, online double auction markets represent the dominant type of exchange market, and traders' manipulations are highly critical in an online double auction market. Thus a robust mechanism that can prevent traders' manipulations and quickly adapt to market changes is a desirable property for an online double auction market.

1.2 Methodologies

This section presents some basic and major approaches used in online mechanism design and, in particular, three new approaches are proposed.

Without knowing who will arrive or what will happen after a decision has been made, it is very difficult for a mechanism/algorithm to make decisions that satisfy some overall goals such as efficiency (that is, maximising social welfare). To understand the difficulty, let us consider a simple example of ranking a set of numbers in an online fashion; that is, the numbers come one by one and on the arrival of each number a final position has to be assigned to this number without knowing what numbers will come afterwards. It is evident that there is no (online) ranking algorithm that can guarantee that the outcome is a proper ranking. The algorithms designed to solve this kind of online problem are called *online algorithms*.

In general, certain goals achievable in a static trading environment cannot be achieved in a corresponding online case, because of, for example, the uncertainty about traders who have not yet arrived. Therefore, in order to measure the performance of an online algorithm/auction, we need to compare the result of an online auction with the optimal (offline) solution. The optimal solution is the best solution with regard to the goals an auction can achieve, given that all future inputs are known to the auction before it makes any decision; that is, there is no uncertainty about the information required for the decision-making. This kind of performance comparison is known as *competitive analysis* [Borodin and El-Yaniv, 1998].

1.2.1 Proposed Approaches

Apart from the above basic online algorithm design techniques, I propose three additional approaches for (online) double auction design in this thesis.

1. *Reduction*: this approach reduces an online double auction to an online one-sided auction. The main idea of this approach is treating sellers as the same as buyers, to let sellers compete with buyers to gain their commodity back if their valuations are competitive.
2. *Behaviour-based auction design*: the main idea of this approach is to use decisions to control or guide the behaviour of participants. Given this control to an online auction, the auction will be able to predict/control the dynamics of the future and therefore more desirable allocations will be achieved.
3. *Augmentation techniques from graph theory*: the advantage of this approach is that the computational complexity is significantly reduced when compared to the traditional approaches. This approach will be used to design a computationally efficient double auction to compute the optimal offline solution for an online double auction.

1.2.2 Existing Approaches

It is worth mentioning other major approaches that have been used in the literature below.

One of them is the use of accessible prior knowledge of the dynamics. Prior knowledge reduces the complexity of the dynamics to some extent. For example, for one-sided dynamic environments, secretary-problem-based online auctions assume that traders arrive randomly [Buchbinder et al., 2010, Hajiaghayi et al., 2004, Kleinberg, 2005]. Another example of online double auctions, assumed that the valuations of traders are limited to a certain range, Blum *et al.* [2006] proposed a truthful online double auction for efficiency in an adversarial setting. Moreover, given that the length of each trader's active time is no more than some constant, Bredin *et al.* [2007] designed a framework to construct truthful online double auctions from truthful static double auctions, and demonstrated the performance

(with regard to efficiency) through experiments of the auctions provided by the framework in probabilistic settings .

Another approach is online learning. For example, Blum *et al.* [2003] demonstrated that through online learning, new auctions substantially improve upon the performance of previous auctions for this problem. Hajiaghayi *et al.* [2004] used learning to design adaptive limited-supply online auctions based on the secretary problem. Parkes and Duong [2007] and Constantin and Parkes [2009] proposed adaptive online auctions, based on online algorithms for stochastic optimisation, by cancelling allocation decisions violating certain properties, such as truthfulness.

Many other approaches have been used; for example, computational difficulty is used as a barrier against manipulations of participants (e.g. [Conitzer and Sandholm, 2007]). They are not listed exhaustively here.

1.3 Major Contributions

Apart from the proposed approaches, the contributions of this thesis can be grouped into three categories: static, approaching online and online. Overall, all the contributions mainly serve online double auction design. The ‘static’ category attacks the auction design problem in simple static environments. The ‘approaching online’ category contains solutions for complex static environments in which the complexity of the auction design is approaching the design difficulty in the corresponding online environments. The ‘online’ category tackles the design problem in complex dynamic environments.

- Static:
 - **I propose an allocation/matching algorithm, called *maximal matching*, to maximise market liquidity for single-unit double auctions.** Maximal matching not only maximises the transaction size, but also shows computational advantage compared with other similar algorithms from the graph theory literature. This algorithm also serves as one of the main components for the behaviour-based online double auction design in the third category.

- Approaching online:
 - **I design a computationally efficient VCG mechanism for bilateral trading environments with temporal constraints, using augmentation techniques from graph theory.** The key advantage of this auction is that it is computationally faster than the classical solution found in the literature. Apart from serving in computing the optimal (offline) solution of one of the online double auctions proposed in the third category, this mechanism also provides a dedicated solution for real trading environments such as futures exchange.
 - **I show the impossibilities of designing certain desirable mechanisms in a group buying online shopping environment, and also provide promising positive results related to these impossibilities.** In particular, no budget-balanced, individually rational and truthful allocations with a reasonable transaction size are present in this environment. In addition to these negative results, I also show promising positive results with budget-balanced, individually rational and truthful or partially truthful mechanisms. These new insights (impossibilities and possibilities) for the design of double auctions under group buying will help us to discover better solutions for the problem.

- Online:
 - **I propose a reduction framework within which to build online double auctions by reducing them to online one-sided auctions.** Well-studied online one-sided auctions can be easily reused in the reduction framework. More importantly, the properties of the reduced online double auction match those of the online one-sided auction, which is difficult to achieve even for static double auctions. In addition, the impossibility of designing a deterministic online double auction that is truthful, individually rational and competitive for efficiency is proved.
 - **I develop a behaviour-based online double auction design approach for a more complicated dynamic environment that simulates real bilateral trading environments.** The novelty of this approach is that it learns traders' behaviour models to guide market decisions so that the auction is able to control/predict the dynamics. To that end, the proposed auction will be able to attract more good traders and make more efficient allocations, which is the hallmark of the success of a real exchange market. The advantage of this approach has been demonstrated through the world-wide market design competitions.

1.4 Related Work

Both one-sided auctions and double auctions involve three different roles: supply (seller(s)), demand (buyer(s)) and the auctioneer (or the market owner). Depending on the causes of the dynamics, online market models are categorised into three different groups here. In the rest of this section, these categories are listed and the main results related to each category.

1.4.1 Static Demand and Dynamic Supply

In this category, the dynamics/uncertainty comes from the supply side. For example, an Ad auction with a fixed number of traders/advertisers, but without knowing how many search queries will come over time [Babaioff et al., 2010, Mahdian and Saberi, 2006].

The model studied by Mahdian and Saberi [2006] is one where there is a fixed number of buyers (advertisers), each of whom bids for the same keyword and only requires one unit, and where the number of incoming searches with that keyword is unknown/dynamic. The goal is to find a single-priced auction to maximise the seller's revenue. The auction is constrained such that each incoming search of the keyword must be assigned to a buyer, otherwise ignored; that is, it cannot be sold in the future. The price for each winning buyer is fixed until all queries have arrived. Mahdian and Saberi proposed two non-deterministic online algorithms. One is 4-competitive¹ to maximise revenue but without considering incentive compatibility, and the other is an extension of the first one with consideration of incentive compatibility, which achieves at least a constant fraction of the optimal revenue. More complicated models of Ad auctions have also been examined, considering multi-unit, multi-keywords or budget constraint (for example, [Mehta et al., 2007]).

1.4.2 Dynamic Demand and Static Supply

On the contrary, in this category the uncertainty is caused by demand. A seller, for instance, has a fixed number of identical items to sell to dynamic arriving and departing buyers; that is, they come over time and also leave the market at different time points.

¹4-competitive means that the online auction achieves a revenue at least $\frac{1}{4}$ of the optimal revenue, where 4 is the *competitive ratio*.

There are two types of this category that are widely studied. One is secretary-problem-based auctions [Babaioff et al., 2008, Hajiaghayi et al., 2004, Kleinberg, 2005], where the supply is limited and also known, and it is the number of items or positions available. The challenge in the classical secretary problem is finding the optimal stopping rule to hire the candidate with the highest quality from n candidates who arrive randomly one after another. The other is unlimited supply auctions, such as digital goods [Bar-Yossef et al., 2002, Blum and Hartline, 2005]. The goal is to find the optimal number of items to sell given that buyers arrive and leave at different time points.

Note that an online auction with reusable goods or pre-scheduled availability of the goods also belongs to this category, which is a special case of secretary-problem-based auctions, for example, [Hajiaghayi, 2005]. Although the supply in this case is not completely fixed or available at the beginning, it is predictable. For example, a fixed number of items is available for sale in each discrete time point $t \in \{1, 2, \dots, T\}$ and they cannot be sold in the future, like ice-cream. Another example of reusable goods is internet bandwidth; if a reusable commodity is available and the auction does not allocate it to anyone, then it is wasted. Thus, the total number of items available for sale is known, though they might not be available all the time.

1.4.3 Dynamic Demand and Dynamic Supply

Market models in this category are the most complex dynamic models with respect to the complexity of the uncertainty. For example, in the dynamic double auctions studied in [Blum et al., 2006, Bredin et al., 2007], the auctioneer is neither a buyer nor a seller, while in all the models introduced above, the auctioneer is either the buyer or the seller. This category is not a simple combination of the above two categories and results from the above two categories cannot be directly applied here. In terms of truthfulness, for example, both buyers and sellers play

strategically in this category, while normally only one side plays strategically in the other two categories.

1.5 Outline of Chapters

Chapter 2 introduces the basic concepts and goals/desiderata of double auction design and the models of both static and dynamic/online double auctions.

Chapter 3 shows the essential allocation/matching algorithms for both static and online double auctions, including the proposed maximal matching.

Chapter 4 initiates the study of online double auctions by adding temporal constraints into trader types without considering the effect of the uncertainty (which is fully examined in Chapters 6 and 7). By utilising augmentation techniques from graph theory, the computational complexity of the classical VCG mechanism is significantly reduced in this model. It shows that the proposed mechanism is $O(n)$ times faster, given n that is the total number of traders.

Chapter 5 tackles the market design problem in the online shopping model with group buying, where traders' valuations are changing with respect to the allocation decisions of the market, in particular, the number of units each trader trade. It provides new insights (impossibilities and possibilities) into the auction design problem under group buying. In particular, it shows that there is no (weakly) budget-balanced, individually rational and truthful mechanism that can also guarantee the transaction size, although there do exist trivial (weakly) budget-balanced, individually rational and truthful mechanisms, for example, the ones with predetermined fixed prices.

Chapter 6 demonstrates how a complicated online double auction can be reduced to an online one-sided auction. More importantly, it shows that the truthfulness and competitiveness (with regard to efficiency) of the reduced online double auction

match/follow those of the online one-sided auction. It also shows that designing a deterministic online double auction that is truthful, individually rational and competitive is impossible, given that no prior knowledge of the uncertainty is accessible. However, such a mechanism is achievable if certain prior knowledge of the dynamics is available.

Chapter 7 studies a more complicated, dynamic bilateral trading environment, where each trader may be active for more than one discrete period of time with a variable valuation in response to the decisions of the mechanism. Since the dynamics are responsive to the decisions of the mechanism, we are able to control or predict the uncertainty to some extent. To that end, this chapter proposes an auction design framework that firstly learns traders' private behaviours and then utilises that learnt result to guide market decisions in order to have better control/prediction of the traders coming in the future so that more desirable allocations will be achieved.

Chapter 8 concludes the major contributions of the thesis and proposes some directions for further investigation.

Chapter 2

Double Auction Models: Static and Dynamic

After a brief introduction to some important double auction design objectives/desiderata, this chapter presents the models for both static and online/dynamic double auction design. Due to the revelation principle [Myerson, 2008b], this thesis will focus on direct-revelation mechanisms. In other words, each trader (seller or buyer) is required to report his or her private information (aka type) directly to the auction. The type report submitted by a seller is known as the **ask**, while the report from a buyer is the **bid**. Both asks and bids are called **orders**.

First, a brief overview of double auction design. We know that a mechanism is a specification of how economic decisions are determined as a function of the information that is known by the individuals in the economy [Myerson, 2008a]. A double auction is one kind of mechanism and its decisions consist of a resource allocation and a payment calculation for each trader. Therefore, to design a double auction is to design an allocation policy and a payment policy. The allocation policy determines who will get the resources/goods, which is a matching between buyer and seller, and the payment policy calculates how much each trader has to

pay. In general, all mechanisms involving payments have the same structure. Although the allocation policy and the payment policy are two separate components, they have to be designed together in order to satisfy certain desiderata that will be introduced next.

2.1 Design Objectives

The following desirable objectives or desiderata are the most commonly considered in double auction design and also in this thesis [Dash et al., 2003, Nisan, 2007]:

- *Incentive Compatibility (or Truthfulness)*. A mechanism is said to be incentive compatible if all of the participants maximise their utilities when they truthfully reveal any private information asked for by the mechanism. This property is also known as truthfulness and truth-telling.
- *Social Welfare Maximisation (or Efficiency)*. This objective corresponds to maximising the goods of the buyers and sellers in aggregate. Specifically, the objective is to have the goods end up in the hands of the agents who value them the most. That is, the goods are allocated to the traders who value them most highly.
- *Budget Balance*. The total payment that the buyers and sellers make equals zero (a strict budget balance), so no money is injected into or removed from the mechanism. We say a mechanism is *weakly budget-balanced* if the total payment is non-negative, so the mechanism does not run at a loss.
- *Individual Rationality*. A mechanism is individually rational if it gives its traders non-negative utility/profit. In other words, the mechanism's allocations do not make any trader worse off than had the trader not participated, so traders volunteer to participate in the mechanism.

- *Profit Maximisation.* Each pair of ask and bid that are matched produces a profit, which is the difference between the bid price and the ask price. This objective is to maximise the sum of these differences, over all matched pairs.
- *Liquidity Maximisation.* The goal is to maximise: (a) the number of transactions, (b) the sell volume, i.e. the total amount of cleared asks, and (c) the buy volume, i.e. the total amount of cleared bids.

2.2 Static Double Auctions

(Static) mechanism design is the discipline of designing mechanisms for static environments. An environment is considered static if all participants' type reports, i.e. the input of the mechanism, are known to the mechanism before the mechanism makes any decision. That is, the mechanism faces no uncertainty about the information it needs for its decision-making.

(Static) double auction design is the design of mechanisms for static bilateral trading environments in which multiple sellers and multiple buyers exchange one commodity simultaneously. The allocation policy of a double auction is a (bipartite) matching between buyers and sellers. A matching is a set of buyer-seller pairs and an exchange quantity is associated with each pair, which indicates the number of units transferred in the pair. It is clear that the exchange quantity is one for all pairs in single-unit environments, where each trader supplies or demands only one unit of a commodity. In the next section, a formal description is provided for the setting and the corresponding definitions of some properties introduced in Section 2.1.

2.2.1 The Model

Consider a double auction market, in which a set B of **buyers** and a set S of **sellers** trade one commodity simultaneously. Buyers and sellers are called **traders**. Let $T = B \cup S$ and assume that traders are independent and $B \cap S = \emptyset^1$; that is, no trader can be both a seller and a buyer.

Each trader $i \in T$ has a privately observed **type** $\theta_i = (v_i)$, $v_i : \mathbb{Z} \rightarrow \mathbb{R}$ is the trader's valuation function, where the input is the number of units of a commodity and the output is the trader's valuation on the bundle of that number of units, e.g. $v_i(3)$ is trader i 's valuation on a bundle of 3 units of the commodity. Let $v(\theta_i) = v_i$. Note that the type θ_i of trader i does not specify explicitly how many units i supplies/demands, but this information can be carried by the valuation function as well, which is denoted by c_i . For instance, if seller i only supplies c_i units as a bundle, then i can set $v_i(\neq c_i) = \infty$, while if buyer j only wants a bundle of c_j units, j can set $v_j(\neq c_j) = 0$, or $v_j(< c_j) = 0$ and $v_j(> c_j) = v_j(c_j)$ if j does not mind getting more than c_j units without extra payment. When discussing single-unit double auctions, the notation is simplified by dropping the input of v_i and using v_i to directly indicate i 's valuation for one or more units.

Let $\theta = (\theta_i)_{i \in T}$ denote the **type profile** where θ_i is the type of trader i . θ_{-i} indicates the type profile of all traders except trader i . Note that a type profile is treated as a vector of types rather than a set of types. Let Θ_i be the set of all possible types of trader i , and let $\Theta = (\Theta_i)_{i \in T}$ be all possible type profiles of all traders in T .

Although traders are required to report their types directly to the auctioneer (that is, the market owner), they do not necessarily report their true types. Let $R(\theta_i)$ be the set of all permitted type reports from trader i of type θ_i , $R(\theta) = (R(\theta_i))_{i \in T}$ be the set of all permitted type profile reports from all traders, $R(\Theta_i) = \bigcup_{\theta_i \in \Theta_i} R(\theta_i)$

¹In the real world, a trader can be both a seller and a buyer for the same commodity. In such a case, it is modelled as two different roles since the decision-making for selling and buying is different.

be the set of all possible reports from i , and $R(\Theta) = (R(\Theta_i))_{i \in T}$ be the set of all possible type profile reports from all traders. Note that in some cases traders misreporting are constrained, so $R(\cdot)$ actually carries those constraints. In this general static model, there is no constraint on misreporting, but examples are shown in this chapter.

Definition 2.1. An **allocation policy** $\pi = (\pi_i)_{i \in T}$ is a function that assigns an integer number to each trader i indicating the number of units traded by i , given traders' type profile report $\hat{\theta} \in R(\Theta)$, such that $\sum_{i \in B} \pi_i(\hat{\theta}) = \sum_{i \in S} \pi_i(\hat{\theta})$.

An allocation policy determines whose order is granted for a transaction and also guarantees that the allocation outcome is **feasible**; that is, the auctioneer never takes a short or long position in the commodity exchanged in the market. For a trader i , if $\pi_i(\hat{\theta}) > 0$ then i **wins**; otherwise i **loses**.

Definition 2.2. A **payment policy** $x = (x_i)_{i \in T}$ is a function that assigns a real number to each trader given traders' type profile report $\hat{\theta} \in R(\Theta)$; that is, $x_i(\hat{\theta}) \in \mathbb{R}$ for all $i \in T$.

Definition 2.3. A **double auction (DA)** on Θ is a pair (π, x) , where π is an allocation policy and x is a payment policy.

Note that the double auction definition above only covers deterministic auctions, and a non-deterministic double auction will be represented as a probability distribution of deterministic double auctions, although sometimes an allocation policy (or a mechanism) is directly defined with probabilistic outcomes.

Given trader i of type $\theta_i = (v_i)$, report profile θ' and DA $\mathcal{M} = (\pi, x)$, the **utility** of i is defined as

$$u(\theta_i, \theta', (\pi, x)) = \begin{cases} v(\theta_i)(\pi_i(\theta')) - x_i(\theta'), & \text{if } i \in B. \\ x_i(\theta') - v(\theta_i)(\pi_i(\theta')), & \text{if } i \in S. \end{cases} \quad (2.1)$$

Considering that DA $\mathcal{M} = (\pi, x)$ might be non-deterministic, $E[u(\theta_i, \theta', (\pi, x))]$ is used to denote the **expected utility** of trader i .

Given the above utility definition, the formal description of the truthfulness property and efficiency property briefly introduced in Section 2.1 is provided. To recall an auction is truthful if reporting type truthfully maximises each trader's utility, and a mechanism is efficient if it always allocates resources to those traders who value them most highly, among all feasible allocations.

Definition 2.4. DA $\mathcal{M} = (\pi, x)$ is **truthful** (aka **incentive-compatible**) if $E[u(\theta_i, (\theta_i, \theta'_{-i}), (\pi, x))] \geq E[u(\theta_i, \theta', (\pi, x))]$ for all $i \in T$, all permitted misreports $\theta' \in R(\theta)$, all type profile $\theta \in R(\Theta)$.

Definition 2.5. DA $\mathcal{M} = (\pi, x)$ is **efficient** if \mathcal{M} maximises the expected **social welfare**

$$E\left[\sum_{i \in B} v(\theta_i)(\pi_i(\theta)) + \sum_{i \in S} v(\theta_i)(c_i - \pi_i(\theta))\right] \quad (2.2)$$

for all type profile $\theta \in R(\Theta)$, where c_i is the number of units seller i supplies.

2.3 Approaching Online Double Auctions

We consider that an environment is dynamic if the individuals in the environment are dynamically arriving and departing, or if their valuations are changing over time, or both. The difference between a static environment and a dynamic environment is that the mechanism for the first faces no uncertainty of the information required for its decision-making, while the mechanism for the latter faces uncertainty. The reasons for this uncertainty in a dynamic environment are that (1) traders are not available at the same time, for example, they are dynamically arriving and departing, and (2) traders' valuations vary over time. These two causes make the information required for decision-making incomplete, and also challenge the decision-making even if there is no uncertainty about the required information.

For example, a seller and a buyer can exchange within a static environment but will not be able to do so in a dynamic environment if they are not available at the same time. Rather than directly going to dynamic environments, this section shows how to model complex static environments such that the decision-making in these environments is as complex as in dynamic environments, except that there is no uncertainty faced in the decision-making. According to the two causes of the dynamics described above, two different complex static environments will be modelled below: one is called *double auction with temporal constraints* and the other is called *double auction under group buying*.

2.3.1 The Model

2.3.1.1 Double Auction with Temporal Constraints

In addition to the valuation function, the type of each trader in this model contains a temporal constraint. The temporal constraint is a period of time. Exchange can occur between two traders if and only if the intersection of their temporal constraints are not empty; that is, the transaction that occurred between two traders without intersection of their temporal constants does not bring any value to either of them. In addition to serving the online double auction design, this model also demonstrates some real static environments. In a futures market, for example, each futures contract is to buy/sell specific quantities of a commodity at a specified price with delivery set at a specified time in the future. A more detailed example would be a seller wanting to sell 100 tonnes of corn in August 2012 at a fixed price, and there is a buyer who wants to buy exactly that amount of corn between June and September 2012 at that price, so they can reach an agreement to exchange 100 tonnes of corn in August 2012.

To model this kind of environment, a period of time is added into traders' type, such that the type of trader i is $\theta_i = (v_i, s_i, e_i)$, where v_i is the valuation function, and s_i and e_i are the starting point and the ending point of the **temporal constraint** $[s_i, e_i]$. To simplify the analysis, we will consider single-unit environments; that is, v_i will be directly used to indicate trader i 's valuation for one or more units.

Given the above extended types, all the concepts described in Section 2.2.1 are still applicable to this model. The only difference here is the allocation policy as it has to consider the temporal constraint to obtain a matching. Detailed matching algorithms to handle this issue are given in Chapters 4 and 6. As seen in Section 2.4, this model is also the corresponding static version of the model for online environments with dynamically arriving and departing traders.

For type reporting, traders do not necessarily truthfully report their types but no early-start and no late-end misreports are permitted. Formally, let $\theta_i = (v_i, s_i, e_i)$ be trader i 's type and $\hat{\theta}_i = (\hat{v}_i, \hat{s}_i, \hat{e}_i) \in R(\theta_i)$ be the trader's report such that $[\hat{s}_i, \hat{e}_i] \subseteq [s_i, e_i]$. The intuition behind this assumption is that no trader would report a temporal constraint that might give that trader negative utility, as traders have no valuation for any transaction that happens outside their temporal constraint.

2.3.1.2 Double Auction under Group Buying

The environment modelled in Section 2.3.1.1 considers temporal constraints that map to the dynamic arrival and departure feature in the corresponding online environment, which is one of the causes of the dynamics described in the very beginning of this section. The other cause of the dynamics is the variation of traders' valuation. To see the variation of traders' valuation in a static environment, a hugely expanding online shopping model, which leverage group buying, is shown. Group buying is a business model in which a number of buyers join

together to order a product in a certain quantity in order to gain a desirable discounted price. Such a business model has recently received significant attention from researchers in economics and computer science, mostly due to its successful application in online businesses, such as *Groupon*.

To model this group buying shopping environment, the model given in Section 2.2.1 is extended by adding extra information/constraints into traders' valuation functions. Assume that sellers' valuation is **monotonic**:

$$v_i(k) \leq v_i(k + 1),$$

and satisfies **group buying discount**:

$$\frac{v_i(k)}{k} \geq \frac{v_i(k + 1)}{k + 1}.$$

That is, a seller's valuation is non-decreasing as the number of units to sell increases, while the mean unit valuation is non-increasing (so buyers can get a discount if the mean valuation is decreasing). One intuition for the group buying discount constraint is that the average unit production cost may decrease when many units can be produced at the same time. For a buyer i of type v_i requiring $c_i > 0$ units, v_i satisfies $v_i(k) = 0$ for all $k < c_i$ and $v_i(k) = v_i(c_i) > 0$ for all $k \geq c_i$. The first constraint of buyers' valuations is that their demands cannot be partially satisfied. The second assumption is that there is no extra value for buyers to get extra units and also no cost for them to dispose of extra units (aka *free disposal*).

Special Challenges

The model in Section 2.3.1.1 can reuse all the formal definitions given in Section 2.2.1. Unfortunately, we cannot reuse all of the notions for this model. The main reason is that the definition of utility in this model is not the same as in

the other models, mainly for sellers. The utility defined in Section 2.2.1 is the most commonly used one, called quasilinear utility, which assumes either that the valuation for each unit is the same if traders' supplies and demands can be partially satisfied, or that they do not have value for partial satisfaction. However, in this model, sellers are encouraged to sell a portion of their supply with different prices, so the valuation for each unit is not the same before and after the auction. Therefore, how to calculate their utilities becomes a difficult question. For instance, a seller supplies two units of a commodity with unit prices $p_1 > p_2$ for selling one and two units, respectively. If one unit is left for the seller, what will be the seller's valuation for this unsold unit? If the valuation for the unsold unit is different before and after the auction, how should it be different and, therefore, how should the seller calculate utility? Thus, the quasilinear utility definition given in (2.1) cannot be used here without additional constraints. This is a very interesting problem in economics, and we are not yet aware of sufficient solutions. This question is left for future work, and in this thesis it is assumed that the supply of each seller is unlimited. Given this unlimited supply assumption, all the notations given in Section 2.2.1 are applicable here.

2.4 Online Double Auctions

An environment is dynamic if its participants are arriving and departing over time [Parkes, 2007], the valuation of each participant is changing over time [Bergemann and Välimäki, 2010, Cavallo and Parkes, 2008], or both [Cavallo et al., 2009]. The mechanism design problem for dynamic settings is termed *online mechanism design*. The main challenge in online mechanism design is that decisions of an online mechanism have to be made dynamically, without knowledge of future participants and/or types. For instance, a seller is selling a house, and each buyer comes at a different time, with a price to buy the house and a waiting period

within which the seller has to decide whether or not to sell it to that buyer. The challenge for the seller is deciding when and to whom to sell the house.

In this thesis, two dynamic environments are examined. In one, traders are arriving and departing dynamically and they only participate in the market for one period of time, with an invariant valuation during their participation. In the other environment, each trader may actively participate in the market for multiple discrete periods of time, and the valuations are not necessarily the same for two different active periods. More importantly, in the first environment, the traders' valuations and active time do not change in response to the decisions of the auction, while the traders' valuations and active time vary in response to the decisions of the auction in the second environment. The first environment is called a *decision-independent dynamic environment* and the second a *decision-dependent dynamic environment*. The first environment is a direct extension of the environment modelled in Section 2.3.1.1. The second is an extension of the first environment plus the extension of the environment modelled in Section 2.3.1.2.

2.4.1 The Model

As these environments are extensions of the environments modelled above, it is possible to reuse most of the formalisations from the above with some minor variations. In the rest of this section, these variations are briefly introduced.

2.4.1.1 Decision-independent Dynamic Environment

To model this environment, the formalisation from Section 2.3.1.1 is reused. The only difference here is the allocation policy, which faces uncertainty for its decision-making in this environment. To distinguish these two different environments, $\theta_i = (v_i, a_i, d_i)$ is used to represent the type of trader i , where $v_i, a_i, d_i \in \mathbb{R}^+$, v_i is

i 's valuation function, and a_i and d_i are the starting point and the ending point of i 's **active time**; that is, the arrival and departure time of i .

Regarding misreports, assume that traders can report any type but no early-arrival and no late-departure misreports are permitted; that is, given trader i 's type $\theta_i = (v_i, a_i, d_i)$, his report $\theta'_i = (v'_i, a'_i, d'_i)$ satisfies $a'_i \leq d'_i$ and $[a'_i, d'_i] \subseteq [a_i, d_i]$. The intuition behind this constraint is that traders do not recognise the market before their arrival and they do not obtain utility for any trade occurring after their true departure time.

2.4.1.2 Decision-dependent Dynamic Environment

This is a highly complex dynamic environment, in the sense that both the trader's valuation and active time are changing over time with respect to the decisions of the auction. To model this kind of environment, *state transition system* has been used, for example, [Cavallo et al., 2009]. Other than model traders' types as any state transition system, this thesis will focus on some special models of traders that are well recognised in real market environments. The auction design issue is investigated in a dynamic environment with these well-studied trader models. The details of these models will be introduced during the design procedure.

2.5 Summary

This chapter has provided an overview of (online) double auction design, especially the design goals. It formally described the models for both the static and dynamic double auction environments studied in this thesis. These environments are grouped in three categories. The first is the general static environment covering both single-unit and multi-unit double auctions studied in the literature. The second category contains two advanced static environments that model certain interesting real environments and also serve the online double auction design in

the environments modelled in the third category. The third category models two dynamic double auction environments which are extensions and combinations of the environments modelled in the first two categories.

The model in the first category will be briefly studied in Chapter 3, those in the second category will be studied extensively in Chapters 4 and 5, and the online models in the third category will be examined in Chapters 6 and 7. These models will be recalled and expanded in the corresponding chapter.

Chapter 3

Matching in a Double Auction

The decisions of a double auction consist of an allocation and a payment calculation for each trader. The allocation is actually a matching between buyers and sellers. A matching is set of buyer-seller pairs, and each pair indicates a transaction between the buyer and the seller. If each trader only supplies or demands one unit of a commodity, then the number of units exchanged in a pair is one, otherwise, the number of units exchanged in a pair needs to be specified explicitly. This chapter introduces some basic matching algorithms for both static and dynamic double auctions, all of them except maximal matching are previously studied in the literature.

3.1 Static Matching

In this section, we introduce two fundamental matching algorithms for single-unit double auctions. One is called *Equilibrium Matching* and the other is called *Maximal Matching*.

3.1.1 Equilibrium Matching

Equilibrium Matching is used to find an equilibrium price p^* which balances the bids and the asks going to be matched so that all the bids with price $p \geq p^*$ and all the asks with price $p \leq p^*$ are matched [Friedman and Rust, 1993]. The algorithm is described below.

Equilibrium Matching

1. Sort all asks (bids) in ascending (descending) order with respect to their price.
2. Based on this sort order, starting at the top, if the ask price is less than or equal to the bid price, add that ask-bid pair to the matching result.
3. Repeat last step until there is no more pair can be added in the matching result.

It is evident that Equilibrium Matching gives an efficient allocation. The equilibrium price is normally determined by the last matchable or the first unmatchable shout pair with respect to the matching order in the algorithm [McAfee, 1992, Wurman et al., 1998].

3.1.1.1 Impossibility

A double auction with Equilibrium Matching or its variants can be incentive compatible or efficient (but not both) with some special payment polices [McAfee, 1992, Wurman et al., 1998]. This impossibility has been shown by Myerson and Satterthwaite [1983] (see Theorem 3.1):

Theorem 3.1. *There does not exist a double auction that is truthful, efficient, individually rational and (weakly) budget-balanced.*

Due to the above impossibility result, we have to sacrifice one of the four objectives/properties in Theorem 3.1. McAfee [1992] proposed a truthful, individually rational and budget-balanced double auction that is not efficient. McAfee's key idea is trade reduction, i.e. reducing the match that gives the least social welfare increase if necessary. McAfee also showed that the proposed auction approaches efficiency if the number of traders approaches infinity. This trade reduction idea has inspired some further work dealing with similar problems in different static exchange environments. For example, Gonen *et al.* [2007] proposed a general trade reduction framework for different exchange environments including multi-unit and combinatorial cases.

Instead of efficiency, other properties have also been extensively considered for sacrifice. The well-known VCG mechanism chooses budget balance [Groves, 1973, Vickrey, 1961]. Wurman *et al.* [1998] proposed single-unit double auctions that are efficient, individually rational, budget-balanced but only partially truthful, i.e. truthful only for either buyers or sellers, and they also showed that there is no multi-unit double auction that is individually rational, efficient, budget-balanced and partially truthful, given that a trader's valuation for each unit is independent of how many units he trades if partial satisfaction is possible or traders do not allow partial satisfaction (which is different from the multi-unit double auction studied in Chapter 5). Under a similar setting to the one studied by [Wurman *et al.*, 1998], Huang *et al.* [2002] proposed multi-unit double auctions that are individually rational, weakly budget-balanced and truthful, but not efficient.

3.1.2 Maximal Matching

The goal of a double auction with Equilibrium Matching or its variants is to produce efficient allocations. However, no double auction mechanism with Equilibrium Matching can maximise liquidity as the uniform clearing price might prohibit some matchable shouts from being matched. In order to maximise the number of matches/transactions, it is essential to allow different matches to be cleared at different prices¹. Otherwise, some matches might be cleared at a price that is not between the ask price and the bid price, i.e. it will act against individual rationality, which is another important desideratum of double auction mechanism design. Based on this idea, we introduce a new matching algorithm, named *Maximal Matching*, which maximises the number of matches. The algorithm is described below, which recursively calls Equilibrium Matching.

Maximal Matching

1. Given an input of shouts, calculate the matching (the set of matched pairs) with Equilibrium Matching, and mark all the matched shouts as matched and all the other shouts as unmatched.
2. Recursively check how many matches Maximal Matching can achieve if the input shouts were matched asks and unmatched bids.
3. Recursively check how many matches Maximal Matching can achieve if the input shouts were unmatched asks and matched bids.
4. Choose the minimum of the numbers from the last two steps as the extra number of matches Maximal Matching can achieve.

¹Sales of identical goods or services that are transacted at different prices is named *price discrimination* [Nagle and Holden, 2001]. In reality, discriminatory pricing may apply to differences in product quality. For example, airlines often offer multiple classes of seats on flights, such as first class and economy class. This is a way to differentiate consumers based on preference, and therefore allows the airline to capture more producer's surplus.

5. Cross match extra matchable shouts with the matched shouts in step 1: the ask in the first matched pair is rematched with the last extra matchable bid, while the bid in the pair is rematched with the last matchable ask, then the second matched pair with the second last extra matchable ask and bid, and so on until all extra matchable shouts are matched.




FIGURE 3.1: Equilibrium Matching vs Maximal Matching

Figure 3.1 shows a matching example of both Equilibrium Matching and Maximal Matching with the same set of shouts, where the numbers are the prices of shouts (other information is omitted), M indicates the last matchable pair with Equilibrium Matching, and the arrowed lines link each matched pair. We can see that Maximal Matching achieves two more matches than Equilibrium Matching does. In addition, Maximal Matching shows computational advantage compared with similar bipartite matching algorithms from graph theory. The corresponding complexity analysis is given in the following.

3.1.2.1 Complexity Analysis

Maximal Matching is equivalent to finding a *maximum bipartite matching* in a bipartite graph $G = (V = (X^{ask}, X^{bid}), E)$, where E contains only one edge for each pair of ask a and bid b if $p(a) \leq p(b)$. Let $n_a = |X^{ask}|$, $n_b = |X^{bid}|$, and n_{em} and n_{mm} be the numbers of matches obtained with Equilibrium Matching and Maximal Matching, respectively. Maximal Matching runs in $O(n_a \log n_a) +$

$O(n_b \log n_b) + O((n_{em})^2)$ time in the worst case, where $O(n_a \log n_a)$ and $O(n_b \log n_b)$ are the complexities of sorting asks and bids (e.g. merge sort), and $O((n_{em})^2)$ is that of the rest of Maximal Matching. The worst case condition for Maximal Matching is that $n_{em} = \min(n_a, n_b) - 1$ holds for all Equilibrium Matchings in Maximal Matching, unless $\min(n_a, n_b) \leq 1$. So we can rewrite the complexity of Maximal Matching as $O(\max(n_a, n_b) \log \max(n_a, n_b) + \min(n_a, n_b)^2)$. As reference, the best known worst-case performance bipartite matching algorithm is *the Hopcroft-Karp algorithm*, which runs in $O(|E| \sqrt{n_a + n_b})$, where $|E| \geq (n_{em})^2$ in our model, time in the worst case [Hopcroft and Karp, 1971]. The worst case condition for the Hopcroft-Karp algorithm in our model is that $n_{em} = \min(n_a, n_b)$, so $O(|E| \sqrt{n_a + n_b}) \geq O(\min(n_a, n_b)^2 \sqrt{n_a + n_b})$. Thus Maximal Matching will outperform the Hopcroft-Karp algorithm in the worst case in our model if the number of shouts are big enough.

3.2 Online Matching

We have seen two basic matching/allocation algorithms for static double auctions in the last section. In this section, I will introduce two basic online matching algorithms for online double auctions from the literature. One is called *Ranking*, which has been well-studied in graph theory for online bipartite matching (where the dynamics comes from only one side) [Karp et al., 1990]. The other one is a very greedy online matching for the case where both sides are dynamic, called *Greedy* [Blum et al., 2006].

3.2.1 The Ranking Algorithm (for One-sided Dynamics)

Karp *et al.* [1990] introduced online bipartite matching, which was one of the first problems to receive the attention of competitive analysis. The input to the problem is a bipartite graph $G = (U, V, E)$, in which the vertices in U arrive in an

online fashion and the edges incident to each vertex $u \in U$ are revealed when u arrives. When a vertex in U arrives, the algorithm may match u to a previously unmatched adjacent vertex in V , if there is one. Such a decision, once made, is irrevocable. The objective is to maximise the size of the matching. Karp *et al.* proposed a matching algorithm called Ranking described below. Regarding the performance of the Ranking algorithm, they demonstrated the following theorem with respect to the matching size. Recall that, to measure the performance of an online algorithm, we compare its results with the optimal offline result (see Section 1.2). The competitive ratio is the minimum of $\{\textit{the optimal result}\}$ divided by $\{\textit{the result of the online algorithm}\}$ for all different instances. For example, Theorem 3.2 says that the Ranking algorithm can achieve an allocation with a matching size at least $1 - 1/e \approx 0.63$ of the optimal matching size.

Theorem 3.2. *The Ranking algorithm achieves a competitive ratio of $e/(e - 1)$ with respect to maximising the size of the matching.*

The Ranking Algorithm

Initialization:

- Choose a random permutation (ranking) σ of the vertices of V .

Online Matching:

Upon arrival of vertex u of U :

- Let $N(u)$ be the set of neighbours of u that have not been matched yet.
- If $N(u) \neq \emptyset$, match u to the vertex $v \in N(u)$ that minimises $\sigma(v)$.

Online Weighted Matching

It is worth mentioning that online weighted matching, as another important online bipartite matching, considers the weight rather than the size of the matching in metric spaces [Kalyanasundaram and Pruhs, 1993]. The weight of a matched pair is the distance between the two vertices. However, the result in online weighted matching cannot be directly applied in online double auctions, because of their special setting. Online weighted matching assumes that the distance between any two vertices is non-negative, while in a general double auction, the weight might be negative if we consider weight as the social welfare change of a matched pair of buyer and seller. Moreover, we need to consider social welfare in online double auctions while the weight considered in online weighted matching is just the change in social welfare. Most importantly, when we consider truthfulness and other properties, online double auctions become more complex as only a perfect matching of online matching and payment calculation is able to satisfy these properties.

3.2.2 The Greedy Algorithm (for Two-sided Dynamics)

The Greedy algorithm is proposed for online double auctions where both buyers and sellers arrive and depart dynamically, which has been extensively studied by Blum *et al.* [2006]. The idea of the algorithm is that on the arrival of each bid (ask), check if it can be matched, if so match it with any ask (bid) that can be matched with the bid (ask). To check the performance of the algorithm, [Blum *et al.*, 2006] showed the following theorem with respect to the matching size, i.e. the matching size is at least half of the optimal one.

Theorem 3.3. *The Greedy algorithm achieves a competitive ratio of 2 for maximizing matching size.*

Note that the impossibility result given in Theorem 3.1 can be easily extended to online double auctions. Of course, to achieve better online performance, we

need more dedicated online algorithms matched with carefully designed payment policies as we will see in the following chapters.

3.3 Summary

We have introduced some basic matching/allocation algorithms for both static and online double auctions. They are very basic, but sufficient to demonstrate the difficulty of the allocation problem and how it can be solved. In the following chapters, we combine allocation and payment together to see the design problem in the different environments introduced in Chapter 2.

Chapter 4

Double Auction with Temporal Constraints

This chapter starts the study of a decision-independent dynamic environment in which traders are dynamically arriving and departing, each trader is active for only one period of time and traders' valuations do not change during their active time. The chapter examines the corresponding static design problem for this environment, and the online auction design problem is addressed in Chapter 6 by using reduction.

It is found that the static allocation problem in this model can be effectively transformed into a weighted bipartite matching in graph theory. The allocation policy is efficient if and only if it corresponds to a maximum-weighted bipartite matching. By using the augmentation technique, this chapter proposes a VCG mechanism in this model and demonstrates the computational advantage of the payment compared with the classical VCG payment (the Clarke pivot payment). The algorithms for both allocation and payment calculation run in polynomial time. It is expected that the method and results provided in this chapter can be applied to the design and analysis of dynamic double auctions. The result is also a

dedicated solution for real trading environments such as futures markets in which temporal information is part of traders' orders.

4.1 Introduction

Although price is the major concern of market clearing in most double auction markets, other factors, such as quantity, quality and temporal constraints, are equally important in some market situations. For example, a futures contract normally specifies not only the price of the underlying commodity but also quantity, quality and settlement date. Nevertheless, most real-world exchange markets are purely price-driven and most studies on double auctions are limited to a single-valued domain [McAfee, 1992, Wilson, 1985]. One reason is that some factors, e.g. quantity and quality, can be eliminated by standardising exchange commodities. However, those attributes with a continuous range or large number of varieties, are hard to standardise.

This chapter considers an extension of the single-valued double auction model that allows traders to specify temporal constraints in their orders. Roughly speaking, an order is written in the form (p, t', t'') , where p stands for the order price and $[t', t'']$ represents the time period when the commodity can be exchanged (not for the order itself). In this extension, a bid and an ask is matchable if and only if the bid price is no lower than the ask price and the intersection of their time constraints is non-empty. We show that the market clearing problem under this extension can be transformed into a weighted bipartite matching. This allows us to use some standard techniques in graph theory, such as augmentation, for the design and analysis of the mechanisms. We prove that an allocation for the double auction is efficient if and only if it corresponds to a maximum weighted bipartite matching of the graph encoding the incoming orders. Based on that, we develop an efficient and incentive-compatible double auction mechanism, i.e. a VCG mechanism [Groves, 1973]. Remarkably, the proposed payment can be implemented much faster than

the classical VCG payment, known as *Clarke pivot payment*, while resulting in the same payments, because it directly uses the abridging and replacing paths generated during the allocation process rather than rerun the allocation algorithm as the Clarke pivot payment does.

It is worth mentioning that, similar temporal information is also used to model the dynamics of a corresponding dynamic environment [Blum et al., 2006, Bredin et al., 2007]. Although the meaning of the temporal information of a trader's type in the online setting (see Chapter 6) is different from that in this setting, a trader's type is modelled in the same way in both settings. Therefore, the mechanism in this model also provides an optimal (offline) solution for a corresponding dynamic market. Such an optimal solution can be directly used for calculating the competitive ratio of an online market-clearing algorithm. Moreover, although orders arrive dynamically, the alternating paths are relatively stable and therefore can be used, for example, to identify potential good orders to find more efficient allocations in an online setting.

This chapter is organised as follows. Section 4.2 briefly introduces the market model and related concepts. In Section 4.3, we introduce a graphic representation for market situations and transfer the market clearing problem into a weighted bipartite matching. Section 4.4 concentrates on the design of an allocation algorithm and a payment algorithm, and proves their desirable properties. A short conclusion is given in Section 4.5 with a brief discussion for future work.

4.2 The Model

Some concepts from Chapter 2 are recalled and expanded here to make this chapter easy to follow. Consider a double auction market, in which a set B of **buyers** and a set S of **sellers** trade one commodity simultaneously. Buyers and sellers are **traders**. Let $T = B \cup S$ and assume that the traders are independent and

$B \cap S = \emptyset$. We also assume that each seller and each buyer supplies and demands a single unit of the commodity.

Each trader $i \in T$ has a privately observed **type** $\theta_i = (v_i, s_i, e_i)$, where v_i , s_i and e_i are non-negative real numbers, v_i is the trader's valuation of a single unit of the commodity, and s_i and e_i are the starting point and the ending point of the time constraint $[s_i, e_i]$. If trader i is a buyer, i obtains **utility** $v_i - p$ if i receives a unit of the commodity within the time interval $[s_i, e_i]$ and pays p ; i obtains zero utility if i pays nothing and does not receive the commodity within the time period. Similarly, if i is a seller, i obtains utility $p - v_i$ if i successfully sells a unit of the commodity within the time period $[s_i, e_i]$ and receives payment p ; i obtains zero utility if i fails to sell the commodity within the time period and no payment is made.

Let $\theta = (\theta_i)_{i \in T}$ denote the type profile where θ_i is the type of trader i . θ_{-i} means the type profile of all traders except trader i . Note that we treat a type profile as a vector of types rather than a set of types. Let Θ_i be the set of all possible types of trader i , and we write $\Theta = (\Theta_i)_{i \in T}$.

Due to revelation principle [Myerson, 2008b], this chapter focuses on direct-revelation mechanisms; that is, traders report their types directly to the auctioneer. However, traders do not necessarily truthfully report their types, but no early-start and no late-end misreports are permitted. Formally, let $\theta_i = (v_i, s_i, e_i)$ be trader i 's type and $\hat{\theta}_i = (\hat{v}_i, \hat{s}_i, \hat{e}_i)$ be the trader's **report**. We assume that $[\hat{s}_i, \hat{e}_i] \subseteq [s_i, e_i]$. The intuition behind the assumption is that no trader would report a temporal constraint that might give him negative utility. Let $R(\theta_i)$ be the set of all permitted reports from trader i given his type θ_i , $R(\Theta_i) = \bigcup_{\theta_i \in \Theta_i} R(\theta_i)$ be the set of all possible reports from i , and $R(\Theta) = (R(\Theta_i))_{i \in T}$.

Given traders' reports $\theta \in R(\Theta)$, an **ask** $\theta_i = (v_i, s_i, e_i)$ (means $i \in S$) and a **bid** $\theta_j = (v_j, s_j, e_j)$ (means $i \in B$) are **matchable** if and only if $v_i \leq v_j$ and

$[s_i, e_i] \cap [s_j, e_j] \neq \emptyset$. That is, the bid's valuation is no less than the ask's valuation, and the intersection of their time constraints is not empty.

An **allocation policy** $\pi = (\pi_i)_{i \in T}$ is a function that assigns 0 or 1 to each trader, given traders' reports $\hat{\theta} \in R(\Theta)$. For a trader i , if $\pi_i(\hat{\theta}) = 1$ we say i wins; otherwise i loses. An allocation determines whose order is granted for a transaction.

A **payment policy** $x = (x_i)_{i \in T}$ is a function that assigns a real number to each trader given an input of traders' reports $\hat{\theta} \in R(\Theta)$, i.e. $x_i(\hat{\theta}) \in \mathbb{R}$ for all $i \in T$.

Definition 4.1. A **double auction mechanism** on Θ is a pair (π, x) , where π is an allocation policy and x is a payment policy.

Following the standard definition, we say that an auction mechanism (π, x) is **efficient** if π maximizes

$$\sum_{i \in B \& \pi_i(\theta)=1} v_i + \sum_{i \in S \& \pi_i(\theta)=0} v_i.$$

for any type profile $\theta = ((v_i, s_i, e_i))_{i \in T}$.

We say that an auction mechanism is **incentive-compatible**, i.e. truthful, if for each trader, reporting his true type is his dominant strategy.

There are a number of alternatives to characterise truthfulness in an auction mechanism. We will use one of them in this chapter based on [Nisan, 2007, Parkes, 2007]. To describe it, we need the following two auxiliary concepts [Parkes, 2007].

For each trader i , we define a partial order \preceq_i on $R(\Theta_i)$:

$$\hat{\theta}'_i \preceq_i \hat{\theta}''_i \text{ iff } \begin{cases} v'_i \geq v''_i \ \& \ [s'_i, e'_i] \subseteq [s''_i, e''_i], & \text{if } i \in S \\ v'_i \leq v''_i \ \& \ [s'_i, e'_i] \subseteq [s''_i, e''_i], & \text{if } i \in B \end{cases}$$

where $\hat{\theta}'_i = (v'_i, s'_i, e'_i)$ and $\hat{\theta}''_i = (v''_i, s''_i, e''_i) \in R(\Theta_i)$.

We say that an allocation policy π is **monotonic** if, for each $i \in T$, $\pi_i(\hat{\theta}_i, \hat{\theta}_{-i}) = 1$ implies $\pi_i(\hat{\theta}'_i, \hat{\theta}_{-i}) = 1$ whenever $\hat{\theta}_i \preceq_i \hat{\theta}'_i$.

Definition 4.2. Given a monotonic policy π and traders' reports $\hat{\theta} \in R(\Theta)$, the **critical value** of trader i of type $\theta_i = (v_i, s_i, e_i)$ is defined as

$$c(\theta_i, \hat{\theta}_{-i}) = \begin{cases} \sup\{v'_i : (v'_i, s_i, e_i) \in R(\theta_i) \wedge \pi_i((v'_i, s_i, e_i), \hat{\theta}_{-i}) = 1\}, & \text{if } i \in S \\ \inf\{v'_i : (v'_i, s_i, e_i) \in R(\theta_i) \wedge \pi_i((v'_i, s_i, e_i), \hat{\theta}_{-i}) = 1\}, & \text{if } i \in B \end{cases}$$

It is undefined if no such v'_i exists.

Now we are ready to describe a characterisation of truthfulness, which will be used in Section 4.4. Theorem 4.3 is based on Theorem 9.36 in [Nisan, 2007] for a single-valued domain and on [Parkes, 2007] for a single-valued online domain. The proof is omitted here as it is similar to the above mentioned theorems.

Theorem 4.3. *A double auction mechanism (π, x) is incentive-compatible if and only if:*

- π is monotonic.
- every winning seller (buyer) is paid (pays) his critical value, and the payment is 0 for losing traders.

4.3 Graph Representation

As assumed in the previous section, each trader has only one unit of a commodity to sell or buy. Transaction must be made in pairs: a seller can only sell his good to a unique buyer, assuming their orders are matchable. This means that to allocate the goods in a double auction is to find matchings between buy orders and sell orders. In such a case we can transform the allocation problem into a matching problem in graph theory. As a result an *efficient* allocation corresponds

to a maximum-weighted bipartite matching. We will first review some concepts related to bipartite matching [West, 2000], encode incoming orders in a bipartite graph, and then show some special properties related to the encoding.

Definition 4.4. A graph $G = (V, E)$ is a **bipartite graph** if the vertex set V consists of two disjoint subsets X and Y , and no edge has both end points in the same subset. For explicitness, we denote the graph as $G = ((X, Y), E)$.

Definition 4.5. Given a traders' report $\theta \in R(\Theta)$, we call $\mathbf{G}_\theta = ((\mathbf{S}^\theta, \mathbf{B}^\theta), \mathbf{E})$ a **bipartite graph generated from θ** if

- $S^\theta = \{\theta_i : i \in S\}$ and $B^\theta = \{\theta_i : i \in B\}$,
- $E = \{(\theta_i, \theta_j) : \theta_i \text{ and } \theta_j \text{ is matchable}\}$.

Definition 4.6. Given a graph G , a **matching** M in G is a set of pair-wise non-adjacent edges, i.e. no two edges share a common vertex. The **size** of M is denoted by $|M|$. A vertex is **matched** if it is incident to an edge in the matching. Otherwise the vertex is **free**.

Given a matching M ,

- an **M -alternating path** is a path in which the edges belong alternatively to M and not to M .
- an **M -augmenting path** is an M -alternating path whose endpoints are free.
- an **M -abridging path** is an M -alternating path whose first edge and last edge are in M .
- an **M -replacement path** is an M -alternating path where one of the endpoints is free and one of the ending edges is in M .

A path is **simple** if it has no repeated vertices. In the rest of this chapter, we will only consider simple paths.

Figure 4.1 shows an example of bipartite representation of eight different type reports. Lines and dashed lines indicate matched edges and free edges respectively, and dots and circles denote matched vertices and free vertices respectively. The value beside each vertex is its valuation. Temporal information is not shown in the graph. Path $(3, 10, 2, 9)$ is an augmenting path, path $(2, 10, 4, 7)$ is an abridging path, and path $(2, 10, 4, 7, 5)$ is a replacement path.




FIGURE 4.1: Example of Alternating Paths

Given a matching M , we can use an M -augmenting path p to augment M by changing all matched edges in p to be free and all the free edges to be matched. By contrast an M -abridging path can be used in the same way to abridge M . Consequently, $|M|$ will increase (decrease) by one with one augmenting (abridging) process. An M -replacement path can be used to replace a bid or an ask in M without changing the status of the other vertices.

Definition 4.7. An allocation policy π is **feasible** if for any traders' reports $\theta \in R(\Theta)$, there is a matching M in the bipartite graph generated from θ such that M exactly covers $\{\theta_i : \pi_i(\theta) = 1\}$.

It follows that any matching in a bipartite graph generated from traders' reports uniquely determines a feasible allocation. In the rest of this chapter, we will only consider feasible allocation policies.

Definition 4.8. Given bipartite graph G_θ , an edge e between $\theta_i = (v_i, s_i, e_i)$ and $\theta_j = (v_j, s_j, e_j)$, where $i \in S$ and $j \in B$, we define the **weight** of e as $w(e) = v_j - v_i$. For any set of edges $E' \subseteq E$, the weight of E' is defined as

$$w(E') = \sum_{e \in E'} w(e).$$

The **weight-increase** of an M -alternating path p is the total weight of free edges in p minus that of matched edges in p :

$$\Delta(p) = w(P - M) - w(P \cap M),$$

where P is the set of all edges in p .

If p is an M -augmenting, M -abridging, or M -replacement path, then $\Delta(p)$ is the net change in the weight of the matching after augmenting, abridging, or replacing by p :

$w(M \oplus p) = w(M) + \Delta(p)$, where $M \oplus p \equiv M \oplus P$, P is the set of all edges in p , and \oplus is the **symmetric difference** operator on sets: $A \oplus B = (A \cup B) \setminus (A \cap B)$.

Lemma 4.9. *Given G_θ , a matching M in G_θ , and an M -alternating path p , $\Delta(p)$ is equal to*

- *the valuation of the bid minus that of the ask of the endpoints of p , if p is an augmenting path.*
- *the valuation of the ask minus that of the bid of the endpoints of p , if p is an abridging path.*
- *the valuation of the free (matched) endpoint minus that of the matched (free) endpoint of p when the endpoints are bids (asks), if p is a replacement path.*

Proof. Without loss of generality, assume p is an augmenting path. Since the weight of each edge is the valuation difference between the incident vertices, the weight of all matched edges in p is the sum of the valuations of all matched bids in p minus that of all matched asks in p , while the weight of all free edges in p is the sum of the valuations of all bids in p minus that of all asks in p . So their difference is the valuation of the free bid minus that of the free ask in p . \square

4.4 Efficient and Truthful Policy Design

In order to design a double auction that is both efficient and truthful, by Theorem 4.3, we need to find an efficient and monotonic allocation policy, and a payment policy that calculates the critical value of each winning trader. Inspired by the similarity between this allocation problem and the weighted matching in a bipartite graph, we first transform the model into a bipartite graph. Within this graph, we show how to efficiently use the well established methods from bipartite matching in the allocation policy, and how to calculate critical payment without running the allocation algorithm again.

4.4.1 Efficient & Monotonic Allocation Policy

With the above graph encoding of traders' reports, we designed an *efficient* allocation policy by adopting the *maximum-weighted bipartite matching* that constructs a maximum-weighted matching by beginning with the empty matching and repeatedly performing augmentations using augmenting paths of maximum weight-increase until a maximum-weighted matching is achieved [Kozen, 1991, Tarjan, 1983]. The resulting allocation policy is called **Maximum-weighted Bipartite Matching Allocation** (*MBM Allocation*), which seeks an allocation that maximises social welfare of any reports θ , by first representing θ in a bipartite graph G_θ , and then applying modified *maximum-weighted bipartite matching* to get a maximum-weighted matching M which determines all winning reports.

We added a more detailed path selection rule in the *maximum-weighted bipartite matching* in order to achieve the *monotonicity* property. The rule is based on the order \preceq_p defined for augmenting paths. Let a sequence of vertices $\theta_1 \circ \dots \circ \theta_n$ denote an augmenting path of length n , which starts from ask θ_1 and ends in bid θ_n . We define \preceq_p on all augmenting paths based on their endpoints:

$$\theta_1 \circ \dots \circ \theta_n \preceq_p \theta'_1 \circ \dots \circ \theta'_m \text{ iff } (v'_1, v_n, s'_1, e_1, s'_m, e_n) \leq_s (v_1, v'_m, s_1, e'_1, s_n, e'_m),$$

where \leq_s is the lexicographic order of two equal length sequences of real numbers: $(r_1^1, \dots, r_n^1) \leq_s (r_1^2, \dots, r_n^2)$ iff $\exists_{1 \leq j \leq n} (r_j^1 \leq r_j^2 \wedge \forall_{1 \leq k < j} (r_k^1 = r_k^2))$. We will use \preceq_p in *MBM Allocation* to distinguish augmenting paths that have the same weight-increase.

Maximum-weighted Bipartite Matching Allocation:

Initialization:

- Encode reports θ in bipartite graph G_θ .
- Set the result matching $M = \emptyset$ for G_θ .

Recursion:

- $AugPath = \{p : \Delta(p) > 0 \text{ and } p \text{ is an } M\text{-augmenting path}\}$.
- $MaxAugPath = \arg \max_{p \in AugPath} \Delta(p)$.
- If $MaxAugPath = \emptyset$, stop recursion.
- Otherwise, let $\hat{p} \in MaxAugPath$ s.t. $p \preceq_p \hat{p}$ for any $p \in MaxAugPath$, and $M = M \oplus \hat{p}$.

Output:

- All reports covered by M win and all the rest lose.

Theorem 4.10. *Maximum-weighted Bipartite Matching Allocation is efficient.*

Before proving Theorem 4.10, we show one essential lemma used in the proof. In the rest of this chapter, π denotes *MBM Allocation*.

Lemma 4.11. *Maximum-weighted Bipartite Matching Allocation is efficient if and only if the maximum-weighted bipartite matching maximizes the weight of the matching.*

Proof. The weight of the matching is $\sum_{\pi_i(\theta)=1 \wedge i \in B} v_i - \sum_{\pi_i(\theta)=1 \wedge i \in S} v_i$, which is equal to $\left(\sum_{\pi_i(\theta)=1 \wedge i \in B} v_i + \sum_{\pi_i(\theta)=0 \wedge i \in S} v_i\right) - \sum_{i \in S} v_i$. $\sum_{i \in S} v_i$ is fixed, so if the weight of the matching is maximised, then $\sum_{\pi_i(\theta)=1 \wedge i \in B} v_i + \sum_{\pi_i(\theta)=0 \wedge i \in S} v_i$ is also maximized, and vice versa. \square

In order to prove Theorem 4.10, by Lemma 4.11, we shall prove that the *maximum-weighted bipartite matching* indeed gives a maximum-weighted matching. To do that, we need the two verified properties of the *maximum-weighted bipartite matching* given in Lemmas 4.12 and 4.13 [Tarjan, 1983].

Lemma 4.12. *Given graph G , let M be a matching of size k of maximum weight among all matchings of size k in G . If we augment M by an augmenting path of maximal weight-increase, then we obtain a matching of size $k + 1$ of maximum weight among all matchings of size $k + 1$ in G .*

Lemma 4.13. *The maximum-weighted bipartite matching will augment along augmenting paths of successively non-increasing weight-increase.*

Proof of Theorem 4.10: By Lemma 4.12, the *maximum-weighted bipartite matching* will give a matching M_k of size k of maximum weight among all matchings of size k after k augmentations. By Lemma 4.13, M_k is also maximum-weighted among all matchings of size at most k if the weight-increase at the k -th augmentation is positive. Therefore, the matching the allocation policy gives, until there is no augmenting path of positive weight-increase, is maximum-weighted among all matchings. \square

Theorem 4.14. *Maximum-weighted Bipartite Matching Allocation is monotonic.*

Although we added a specific path selection rule based on \preceq_p to avoid randomisation of *MBM Allocation* in most cases, there is still one situation where \preceq_p cannot help. When two types are the same and two augmenting paths of maximum positive weight-increase start from them and end in the same vertex, then \preceq_p cannot separate these two paths clearly, i.e. both of them have a chance of being selected but none of them are guaranteed. Thus we assume that all type reports of sellers (buyers) are different, i.e. two sellers (buyers) do not share the same type. Note that there might be more than one augmenting path with the same endpoints, but this does not affect the deterministic property of *MBM Allocation*, though it will randomly select one of them to augment.

Proof of Theorem 4.14: By contradiction, without loss of generality, assume that $\pi_i(\theta) = 1$ and $\pi_i(\theta'_i, \theta_{-i}) = 0$ for some bids $\theta_i \preceq_i \theta'_i$. Let θ_i be matched in round k of $\pi(\theta)$, i.e some augmenting path ending with θ_i is of maximal weight-increase in round k . Since θ and θ'_i are both not matched before round k , so the matchings are the same in both $\pi(\theta)$ and $\pi(\theta'_i, \theta_{-i})$ after any round $< k$. Let $\theta_m \circ \dots \circ \theta_i$ be the augmenting path of maximal weight-increase selected in round k of $\pi(\theta)$. Since $\theta_i \preceq_i \theta'_i$, $\theta_m \circ \dots \circ \theta'_i$ is an augmenting path in round k of $\pi(\theta'_i, \theta_{-i})$ and $\theta_m \circ \dots \circ \theta_i \preceq_p \theta_m \circ \dots \circ \theta'_i$. Moreover, in round k , all augmenting paths in $\pi(\theta'_i, \theta_{-i})$, except those that end with θ'_i , are also augmenting paths in $\pi(\theta)$. Thus, in round k of $\pi(\theta'_i, \theta_{-i})$, for any augmenting path p that does not end with θ'_i , $p \preceq_p \theta_m \circ \dots \circ \theta'_i$, and all the rest end with θ'_i . Therefore, an augmenting path ending with θ'_i should be selected in round k of $\pi(\theta'_i, \theta_{-i})$, which contradicts the assumption. \square

4.4.2 Truthful Payment Policy

We have found an efficient allocation policy, *MBM Allocation*, and proved its monotonicity property which is one of the two iff conditions to satisfy truthfulness. What is left is to calculate the critical value for each winning trader.

It is not practical to calculate the critical value as it's defined in Definition 4.2. Here we propose another approach which is inspired by the reverse of *MBM Allocation*. A type θ_i is matched because there is an augmenting path of maximum positive weight-increase ending with θ_i in some round of the matching procedure. Therefore, if a type does not satisfy the above condition, it will not be matched. The basis of our payment policy is to seek the least violation condition for each winning type, i.e. the edge condition between winning and losing.

Given traders' reports θ , if $\pi_i(\theta) = 1$, the payment for trader i , $x_i(\theta)$, is defined in terms of abridging and replacement paths starting from θ_i in the following, which is called **Min-Max Payment** (*MM Payment*). $x_i(\theta) = 0$ if $\pi_i(\theta) = 0$.

Min-Max Payment:

$$x_i(\theta) = \begin{cases} \min_{p \in D \cup R} v(\text{ending}(p)), & \text{if } i \in S \\ \max_{p \in D \cup R} v(\text{ending}(p)), & \text{if } i \in B \end{cases}$$

where

- D is a set of all abridging paths starting from θ_i ,
- R is a set of all replacement paths starting from θ_i ,
- and $v(\text{ending}(p))$ is the valuation of the ending vertex, the endpoint other than θ_i , of path p .

For each winning ask, *MM Payment* gives the minimum valuation such that, if the ask's valuation were greater than or equal to that minimum, it can be removed from the matching to (weakly) increase the weight of the matching, while for each winning bid, the payment is the corresponding maximum. The set D gives all possible ways to remove θ_i by abridging, while the set R gives all possible ways to

substitute a free vertex for θ_i . Note that set D does not necessarily contain the path that was used to match θ_i , as the path can be changed with other augmentations after adding θ_i .

Theorem 4.15. *Given bipartite graph G_θ and a winning type $\theta_i = (v_i, s_i, e_i)$ determined by MBM Allocation, Min-Max Payment $x_i(\theta)$ is equal to critical value $c(\theta_i, \theta_{-i})$.*

To prove Theorem 4.15, we need the following two lemmas which can be found in [Blum et al., 2006, Kozen, 1991].

Lemma 4.16. *Given two matchings M and M' , $M \oplus M'$ consists of a collection of vertex-disjoint alternating paths and even length cycles.*

Proof. By the definition of matching, no vertex can have more than one incident edge from M (or M'), so no vertex can have more than two incident edges from $M \oplus M'$. □

Lemma 4.17. *Given two matchings M and M' , a vertex v is an endpoint of a path in $M \oplus M'$ if and only if it is matched in either M or M' but not both.*

Proof. If vertex v is an endpoint of an alternating path, i.e. there is only one edge incident to v in $M \oplus M'$, then v can only be matched in either M or M' but not both. If vertex v is only matched in only one of M and M' , then v must be contained in $M \oplus M'$ with only one edge incident to it. □

Now we are ready to prove Theorem 4.15.

Proof of Theorem 4.15: Without loss of generality, assume $\theta_i = (v_i, s_i, e_i)$ is a winning ask, and let $x^i = x_i(\theta)$ and $c^i = c(\theta_i, \theta_{-i})$. To prove $x^i = c^i$, by the definition of c^i , we need to show that for any $\theta'_i = (v'_i, s_i, e_i)$:

1. $\forall \theta'_i: \pi_i(\theta'_i, \theta_{-i})=1 (v'_i \leq x^i)$.

2. $\exists_{x^i - \delta < v'_i < x^i} (\pi_i(\theta'_i, \theta_{-i}) = 1)$ for any $\delta > 0$.

Let M be the matching of $\pi(\theta)$ and M' be that of $\pi(\theta'_i, \theta_{-i})$. We will prove these two conditions one by one blow.

Part I: By contradiction, assume that $\pi_i(\theta'_i, \theta_{-i}) = 1$ and $v'_i > x^i$. Let A_M (B_M) be all the matched asks (bids) in M , and $A_{M'}$ ($B_{M'}$) be all the matched asks (bids) in M' . Since π is monotonic and M and M' are maximum-weighted, it follows that all matched asks in M except for θ_i must be matched in M' , i.e. $A_M \setminus \{\theta_i\} \subset A_{M'}$, and all the matched bids in M' must be matched in M , i.e. $B_{M'} \subseteq B_M$. Thus inequalities $|A_M| - 1 < |A_{M'}|$ and $|B_{M'}| \leq |B_M|$ hold. Moreover, $|A_M| = |B_M|$ and $|A_{M'}| = |B_{M'}|$, so we get $|M| = |M'|$, $B_M = B_{M'}$, and $A_M \setminus \{\theta_i\} = A_{M'} \setminus \{\theta'_i\}$. Therefore, by Lemma 4.16 and 4.17, there is only one alternating path $p_{only} = \theta_i \circ \dots \circ \theta'_i$ in $M \oplus M'$, and all the rest are cycles. If all vertices reachable from θ_i through M -abridging or M -replacement paths are also reachable from θ'_i through M' -abridging or M' -replacement paths, then, since $v'_i > x^i$, there is at least one M' -abridging or M' -replacement path of positive weight-increase by which we can remove θ'_i to increase the weight of the matching, which contradicts the choice of M' .

We now prove that the above reachability condition holds. (1). For any vertex v except for θ_i (θ'_i) in p_{only} , the path between θ_i (θ'_i) and v is either an abridging or a replacement path with respect to M (M'). (2). Any vertex v' not in p_{only} that is reachable from θ_i by an abridging or replacement path p is also reachable from θ'_i through the same type of path p' . Since p must be connected with p_{only} and for any edge $e \in p$ and $e \notin p_{only}$, if $e \in M$ and $e \notin M'$, there must be an even length cycle that contains e in $M \oplus M'$, and vice versa, i.e. if e connects vertices v_1 and v_2 in p , there is always a corresponding edge or path connecting v_1 and v_2 in p' . For instance, Figure 4.2 shows one alternating path (a, b, c, d, e) and a cycle (h, i, j, k) of $M \oplus M'$: thin lines and thick lines belong to M and M' respectively, while the double line between f and g is in both matchings and




FIGURE 4.2: Reachability Example

dashed lines are free. It is easy to see that all vertices reachable from a through a M -augmenting or M -replacement path is also reachable from e by a corresponding path with respect to M' .

Part II: To prove the second condition, we will prove $\pi_i(\theta'_i, \theta_{-i}) = 1$ for any $v'_i < x^i$. By contradiction, assume that $v'_i < x^i$ and $\pi_i(\theta'_i, \theta_{-i}) = 0$. By Lemma 4.17 there is a path $p_{\theta_i} \in M \oplus M'$ starting from θ_i and ending with θ_n in either M or M' . Since $x^i \leq v_n$ by the definition of x^i and $v'_i < x^i$ as we assumed, we can substitute θ'_i for θ_i in p_{θ_i} to get an M' -alternating path $p_{\theta'_i}$. If θ_n is matched in M , then $p_{\theta'_i}$ is an M' -augmenting path and by Lemma 4.9 $\Delta(p_{\theta'_i}) = v_n - v'_i > 0$, which contradicts the choice of M' . Thus θ_n is a matched ask in M' , and $p_{\theta'_i}$ is an M' -replacement path. Since M' is a maximum-weighted, by Lemma 4.9 $\Delta(p_{\theta'_i}) = v_n - v'_i \leq 0$. Put all results together, we get contradiction $v_n \leq v'_i < x^i \leq v_n$. \square

Another appealing property of Min-Max payment is its independence from the allocation algorithm. We show that Min-Max payment results in the same payments as the most desirable VCG payment (Clarke pivot payment), but it does not require the recall of the allocation algorithm. Clarke pivot payment is defined as $x_i(\theta) = V^\pi(\theta_{-i}) - V_{-i}^\pi(\theta)$, where $V^\pi(\theta)$ is the social value given traders' report profile θ and the allocation policy π , and $V_{-i}^\pi(\theta)$ is the social value without counting trader i .

Proposition 4.18. *Given traders' report θ and efficient and monotonic allocation policy π , for each trader i , Min-Max payment $x_i^{MM}(\theta)$ is equal to Clarke pivot payment $x_i^C(\theta)$.*

Proof sketch. We need to prove that for each winning type θ_i if we remove θ_i from the maximum-weighted matching M of bipartite graph G_θ by using the path p that gives $x_i^{MM}(\theta)$, the result matching M' is also maximum-weighted in $G_{\theta_{-i}}$. By contradiction, assume that M' is not maximum-weighted, we will conclude either M is not maximum-weighted or path p contradicts the definition of Min-Max payment. \square

Corollary 4.19. *Double auction mechanism (MBM Allocation, MM Payment) is efficient, incentive-compatible and individual-rational, i.e. traders never receive negative utility.*

Figure 4.3 shows an example of the double auction we have defined, where the number beside each vertex is the valuation of the vertex and the value inside parentheses is the payment.




FIGURE 4.3: MBM Allocation and MM Payment

4.4.3 Computational Complexity

We further show that both our allocation policy and payment policy can be implemented in polynomial time and, more importantly, our payment can be implemented much faster than Clarke pivot payment.

Theorem 4.20. *Let n be the number of traders' reports. MBM Allocation can be implemented in time $O(n^3)$, and Min-Max Payment can be implemented in time $O(n^3)$.*

The result given in Theorem 4.20 is significant because, to the best of my knowledge, the implementations of Clarke pivot payment cannot avoid the recall of the allocation algorithm [Nisan et al., 1999, Sandholm, 2003]. In other words, for each winning report θ_i , π needs to search another allocation on the remaining reports θ_{-i} . Therefore, it will take $O(n)$ times of the allocation time in this model, i.e. $O(n^4)$ with *MBM Allocation*.

Proof of Theorem 4.20: Bipartite graph representation of the reports takes at most $n^2/4$ time by checking each pair of ask and bid, so there will be at most $m = n^2/4$ edges. According to [Galil, 1986], finding an augmenting path of maximal weight increase can be solved by Dijkstra's algorithm taking $O(m + n \log n)$ time. There are at most $n/2$ rounds, so *MBM Allocation* can be implemented in time $O(n^3)$. For each winning type, *MM Payment* can be done by depth-first or breadth-first search which takes $O(n + m)$ time. There are at most n winning types, so *MM Payment* can also be implemented in time $O(n^3)$. \square

4.5 Summary

We have developed an efficient and truthful double auction mechanism (i.e. a VCG mechanism) in a model where multiple sellers and multiple buyers exchange a commodity and each trader has a privately observed type consisting of a valuation of a commodity and a time period which constrains when the commodity can be exchanged. The mechanism is characterised by an allocation policy and a payment policy. By encoding the model in a bipartite graph, we efficiently adapted the *maximum-weighted bipartite matching* to get an efficient and monotonic allocation policy. We also developed a truthful payment policy that can be implemented faster than Clarke pivot payment while resulting in the same payments as Clarke pivot payment.

Myerson and Satterthwaite [1983] proved that there is no efficient, incentive-compatible and individual-rational bilateral trade without outside subsidies, i.e. a market with our mechanism will run in deficit. To avoid this deficit, we need to compromise between efficiency and truthfulness. There are two possible remedies: either relaxing efficiency, or giving up incentive compatibility, as investigated by McAfee [1992] and Wurman *et al.* [1998] in single-valued domains. Finding how these compromises can lead to a realistic mechanism under this model is worth further investigation.

Another direction for the future work is to further extend our framework to allow generic constraints on orders. Some real-world markets allow or have a demand for conditional orders. For instance, a trader, in day trading, would send a contingent order, stop order or exit order to his or her broker or the market.

Chapter 5

Multi-unit Double Auction under Group Buying

Chapter 4 has shown the difficulty of the auction design problem for double auction with temporal constraints. This chapter studies the design problem in an environment in which traders' valuations pose the challenge to the design.

This chapter deals with the market situation when multiple sellers sell one kind of product to a number of buyers and a seller is willing to give bigger discount if more buyers group together to buy the product from the seller, which is called group buying. Group buying is a business model in which a number of buyers join together to order a product in a certain quantity in order to gain a desirable discounted price. Such a business model has recently received significant attention from researchers in economics and computer science, mostly due to its successful application in online businesses, such as *Groupon*. We consider this problem as a multi-unit double auction. We first examine two deterministic mechanisms that are budget balanced, individually rational and only one-sided truthful, i.e. it is truthful for either buyers or sellers. Then we find that, although there exists a 'trivial' (non-deterministic) mechanism that is (weakly) budget balanced, individually rational and truthful for both buyers and sellers, such a mechanism is not

achievable if we further require that both the trading size and the payment are neither seller-independent nor buyer-independent. In addition, we show that there is no budget balanced, individually rational and truthful mechanism that can also guarantee transaction size.

Although have seen difficulties for designing certain mechanisms in this model, real environments can be more complicated than that. In real group buying markets, each trader receiving the discounted product might return in the future and also influence their friends to try this product; that is, the advertising effect, which has not been well studied.

5.1 Introduction

Group buying (or collective buying power) is when a group of consumers come together and use the old rule of thumb, there is power in numbers, to leverage group size in exchange for discounts. Led by *Groupon*, the landscape for group buying platforms has been growing tremendously during last few years. Because of the advent of social networks, e.g. *Facebook*, this simple business concept has been leveraged successfully by many internet companies. Taking the most successful group buying platform *Groupon* for example, a group buying deal is carried out in the following steps:

1. the company searches good services and products (locally) that normally are not well-known to (local) consumers,
2. the company negotiates with a target merchant for a discounted price for their services and the minimum number of consumers required to buy their services in order to get this discount,

3. the company promotes the merchant's services with the discounted price and the minimum number of buyers required to make the deal within a period, say two days,
4. if the number of consumers willing to buy the services reaches the minimum during that period, then all the consumers will receive the services with the discounted price, and the company and the merchant will share the revenue. Otherwise, no deal and no loss for any party, especially the merchant and consumers.

All participants benefit from successful group buying deals: consumers enjoy good services with lower prices, merchants promote their services and most likely more consumers will buy their services with normal prices in the future (i.e. group buying also plays the role of advertising), and the company providing the platform benefits from merchants' revenue.

Besides its simple concept and its successful business applications, group buying is not well studied in academia [[Anand and Aron, 2003](#), [Arabshahi, 2011](#), [Byers et al., 2011](#), [Edelman et al., 2011](#)]. In particular, the combination of collective buying power and advertising challenges theoretical analysis. In this work, we extend the simple concept, used by *Groupon*, to allow merchants (or sellers) and consumers (or buyers) to express more of their private information. More specifically, instead of one single discounted price for selling a certain number of units of a product, sellers will be able to express different prices for selling different amounts of the product. Buyers will be able to directly reveal the amount they are willing to pay for a product, other than just show interest in buying a product coming with a fixed price. To that end, we do not just enhance the expression of traders' private information, but also reduce the number of no-deal failures that happen when the number of buyers willing to purchase a product does not reach the predetermined minimum on the *Groupon* platform. Moreover, we will allow multiple sellers to build competition for selling identical products.

Given the above extension, what we get is a multi-unit double auction, where there are multiple sellers and multiple buyers exchanging one commodity and each trader (seller or buyer) supplies or demands multiple units of a commodity. Different from the multi-unit double auctions previously studied [Chu, 2009, Huang et al., 2002], the focus of this model is on group buying and we assume that sellers have unlimited supply and a seller's average unit price is decreasing (non-increasing) when the number of units sold is increasing.

Due to revelation principle, we only consider mechanisms where traders directly report their types. We propose/examine some mechanisms in terms of, especially, *budget balance*, *individual rationality*, and *truthfulness*, which are three important criteria we usually try to achieve in double auction design. Budget balance guarantees that the market owner running the auction does not lose money. Individual rationality incentivises traders to participate in the auction, as they will never get negative utility/benefit for participating in the auction. Truthfulness makes the game much easier for traders to play, because the best strategy can be easily computed for each trader, which is just his true type. Truthfulness also plays an important role for achieving other properties based on traders' truthful types, e.g. *efficiency* (i.e. social welfare maximisation). We will not measure social welfare in this model, because of unlimited supply. However, we will consider the number of units exchanged, called *trading size*, which is part of *market liquidity*, indicating the success of an exchange market.

We find that, even without considering other criteria, budget balance, individual rationality and truthfulness are hard to be satisfied together in this model. We show that there is no budget-balanced, individually rational and truthful auction, given that both the trading size and the payment are neither seller-independent nor buyer-independent. We say a parameter of a mechanism is seller-independent (buyer-independent) if its value does not depend on sellers' (buyers') type reports. However, by allowing either the trading size or the payment to be seller-independent or buyer-independent, we will be able to design auctions satisfying

budget balance, individual rationality and truthfulness at the same time. Even under the non-predetermination constraint, we will propose two mechanisms that are budget-balanced, individually rational and one-sided truthful, i.e. truthful for either buyers or sellers. In addition, we prove that there is no budget-balanced, individually rational and truthful mechanism that can also guarantee trading size.

This chapter is organised as follows. After a brief introduction of the model in Section 5.2, we propose two budget-balanced, individually rational and partially truthful (deterministic) mechanisms in Section 5.3 and 5.4. Following that, we further study the existence of (weakly) budget-balanced, individually rational and truthful mechanisms in Section 5.5. Finally, we conclude in Section 5.6 with related and future work.

5.2 The Model

We study a multi-unit double auction where multiple sellers and multiple buyers exchange one commodity. Each seller supplies an unlimited number of units of a commodity and each buyer requires a certain number of units of the commodity. Each **trader** (seller or buyer) i has a privately observed valuation function (aka **type**) $v_i : \mathbb{Z}^+ \rightarrow \mathbb{R}^+$ where the input of the function is the number of units of the commodity and the output is the valuation for those units together.

We assume that sellers' valuation is **monotonic**: $v_i(k) \leq v_i(k+1)$, and satisfies **group buying discount**: $\frac{v_i(k)}{k} \geq \frac{v_i(k+1)}{k+1}$. That is, a seller's valuation is non-decreasing as the number of units to sell increases, while the mean unit valuation is non-increasing (so buyers can get a discount if the mean valuation is decreasing). One intuition for group buying discount constraint is that the average unit production cost may decrease when many units can be produced at the same time. For a buyer i of type v_i requiring $c_i > 0$ units, v_i satisfies $v_i(k) = 0$ for all $k < c_i$ and $v_i(k) = v_i(c_i) > 0$ for all $k \geq c_i$. The first constraint of buyers' valuation

states that their demands cannot be partially satisfied. The second assumption states that there is no cost for buyers to deal with extra units allocated to them (*free disposal*). Following [Chu, 2009, Huang et al., 2002], we assume that c_i of buyer i is common knowledge. Without loss of generality, we will assume that $c_i = 1$ for each buyer i to simplify the rest of the analysis, and the results under this assumption can be easily extended to the general case.

To participate in an auction, each trader is required to report some information (often related to his type) to the auctioneer (i.e. the market owner). Because of the revelation principle [Myerson, 2008b], we will focus on auctions that require traders to directly report their types. However, traders do not necessarily report their true types.

Let S be the set of all sellers, B be the set of all buyers, and $T = S \cup B$. We assume that $S \cap B = \emptyset$. Let $v = (v_i)_{i \in T}$ denote the type profile of all traders. Let $v_{-i} = (v_1, v_2, \dots, v_{i-1}, v_{i+1}, \dots, v_n)$ be the type profile of all traders except trader i . Given trader i of type v_i , we refer to $R(v_i)$ as the set of all possible type reports of i . Similarly, let $R(v)$ be the set of all possible type profile reports of traders with type profile v . We will use $v^B = (v_i)_{i \in B}$ to denote the type profile of buyers, and $v^S = (v_i)_{i \in S}$ for sellers.

Definition 5.1. A **multi-unit double auction** (MDA) $\mathcal{M} = (\pi, x)$ consists of an **allocation policy** $\pi = (\pi_i)_{i \in T}$ and a **payment policy** $x = (x_i)_{i \in T}$, where, given traders' type profile report v , $\pi_i(v) \in \mathbb{Z}^+$ indicates the number of units that seller (buyer) i sells (receives), and $x_i(v) \in \mathbb{R}^+$ determines the payment paid to or received by trader i .

Note that the above definition of MDA includes only deterministic MDAs, i.e. given a type profile report, the allocation and payment outcomes are deterministic. We will also consider non-deterministic/random MDAs where the outcomes are random variables. A non-deterministic MDA can be described as a probability distribution over deterministic MDAs.

Given an MDA $\mathcal{M} = (\pi, x)$ and type profile v , we say trader i **wins** if $\pi_i(v) > 0$, **loses** otherwise. An allocation π is **feasible** if $\sum_{i \in B} \pi_i(v) = \sum_{i \in S} \pi_i(v)$ and for all S, B and v . An MDA $\mathcal{M} = (\pi, x)$ is feasible if π is feasible. A non-deterministic MDA is feasible if it can be described as a probability distribution over feasible deterministic MDAs. Feasibility guarantees that the auctioneer never takes a short or long position in the commodity exchanged in the market. In the rest of this chapter, only feasible MDAs are discussed.

Given traders' type profile v , their type profile report $\hat{v} \in R(v)$ and deterministic MDA $\mathcal{M} = (\pi, x)$, the **utility** of trader i with type v_i is defined as

$$u(v_i, \hat{v}, (\pi, x)) = \begin{cases} v_i(\pi_i(\hat{v})) - x_i(\hat{v}), & \text{if } i \in B. \\ x_i(\hat{v}) - v_i(\pi_i(\hat{v})), & \text{if } i \in S. \end{cases}$$

Considering \mathcal{M} might be non-deterministic, we use $E[u(v_i, \hat{v}, (\pi, x))]$ to denote the expected utility of trader i .

Definition 5.2. An MDA $\mathcal{M} = (\pi, x)$ is **truthful** (or **incentive-compatible**) if $E[u(v_i, (v_i, \hat{v}_{-i}), (\pi, x))] \geq E[u(v_i, \hat{v}, (\pi, x))]$ for all $i \in T$, all $\hat{v} \in R(v)$, all v .

In other words, a mechanism is truthful if reporting type truthfully maximises each trader's utility. We say an MDA \mathcal{M} is **buyer-truthful** (**seller-truthful**) if \mathcal{M} is truthful for at least all buyers (sellers).

An MDA is **budget-balanced** (BB) if the payment received from buyers is equal to the payment paid to sellers, and it is **weakly budget-balanced** (WBB) if the payment received from buyers is greater than the payment paid to sellers. An MDA is **individually rational** (IR) if it gives its participants non-negative utility. Because of unlimited supply, we will not be able to measure social welfare in this model, as it will be unbounded before and after the auction. *Market liquidity*, as an important indicator of a successful exchange market, will be considered. We will check one of the important measures of market liquidity, the number of units exchanged, called **trading size**.

Given the type profile report v , assume that $v_1^B(1) \geq v_2^B(1) \geq \dots \geq v_m^B(1)$, we define the **optimal trading size** $k_{opt}(v)$ as

$$k_{opt}(v) = \max_k \left(\sum_{i=1}^k v_i^B(1) \geq \min v_j^S(k) \right). \quad (5.1)$$

That is, optimal trading size is the maximal number of units that can be exchanged in a (weakly) budget-balanced auction where the payment of a winning trader is his valuation for receiving/selling the number of units allocated to him. As we will see, it is often not possible to achieve the optimal trading size, if we consider other properties such as truthfulness, individual rationality and budget balance at the same time. Therefore, we define the following notion to measure an MDA's trading size.

Definition 5.3. An MDA \mathcal{M} is **c-competitive** if the (expected) trading size $k_{\mathcal{M}}(v)$ of \mathcal{M} is at least $\frac{k_{opt}(v)}{c}$, for all type profile report v . We say \mathcal{M} is **competitive** if \mathcal{M} is c -competitive for a constant c . We refer to c as competitive ratio.

Moreover, other than following Definition 5.2, we will use Proposition 5.4 to analyse the truthfulness of an MDA. Proposition 5.4 is based on Proposition 9.27 of [Nisan et al., 2007], and its proof directly follows the proof there.

Proposition 5.4 (Proposition 9.27 of [Nisan et al., 2007]). *An MDA $\mathcal{M} = (\pi, x)$ is truthful if and only if it satisfies the following conditions for every trader i with type v_i and every v_{-i}*

- *If $E[\pi_i(v_i, v_{-i})] = E[\pi_i(\hat{v}_i, v_{-i})]$, then $E[x_i(v_i, v_{-i})] = E[x_i(\hat{v}_i, v_{-i})]$. That is, the payment of i does not depend on v_i , but only on the alternative allocation result.*
- *$E[u(v_i, v, (\pi, x))] \geq E[u(v_i, (\hat{v}_i, v_{-i}), (\pi, x))]$ for all $\hat{v}_i \in R(v_i)$. That is, the expected utility of i is optimised by \mathcal{M} .*

5.3 A BB, IR and Buyer-truthful MDA

A Vickrey auction is a truthful and individually rational one-sided auction for exchange of one item, where traders report their private types (valuations for the item), and in which the trader with the highest valuation wins, but the price paid is the second-highest valuation. We apply a similar principle in this section and propose an MDA, called Second Price MDA. We show that this auction is budget-balanced and individually rational but only buyer-truthful.

Second Price MDA \mathcal{M}_{2nd}

Given type profile report $v = (v^B, v^S)$, assume that $v_1^B(1) \geq v_2^B(1) \geq \dots \geq v_m^B(1)$.

1. Let $w(k) = \min \arg \min_i v_i^S(k)$ and $p(k) = \min_{i \neq w(k)} \frac{v_i^S(k)}{k}$ or ∞ if there is only one seller.
2. Let $k^* = \max\{k | v_k^B(1) \geq p(k)\}$.
3. The first k^* buyers, i.e. buyers of valuation $v_1^B, v_2^B, \dots, v_{k^*}^B$, receive one unit of the commodity each and each of them pays $p(k^*)$.
4. Seller $w(k^*)$ sells k^* units of the commodity and receives payment $p(k^*) \cdot k^*$.
5. The remaining traders lose without payment.

Given the number of units going to be exchanged k , \mathcal{M}_{2nd} selects the seller with lowest valuation for selling k units to win (i.e. $w(k)$) and the payment is the second lowest valuation (i.e. $p(k) \cdot k$). k^* of \mathcal{M}_{2nd} , the trading size, is the maximal number of units that can be exchanged, given that each winning buyer pays the mean unit price $p(k^*)$. It is evident that the profit of the auctioneer running \mathcal{M}_{2nd} will be

zero and no participant will get negative utility, i.e. \mathcal{M}_{2nd} is budget-balanced and individually rational.

Lemma 5.5. *For any $k \geq 1$, $p(k)$ of \mathcal{M}_{2nd} satisfies $p(k+1) \leq p(k)$ and $p(k+1) \cdot (k+1) \geq p(k) \cdot k$.*

Proof. Since sellers' valuation satisfies group buying discount, i.e. $\frac{v_i^S(k+1)}{k+1} \leq \frac{v_i^S(k)}{k}$, we get $p(k+1) = \min_{i \neq w(k+1)} \frac{v_i^S(k+1)}{k+1} \leq \min_{i \neq w(k)} \frac{v_i^S(k)}{k} = p(k)$. In other words, the mean unit price is non-increasing as the number of units sold together increases.

Because of $v_i(k+1) \geq v_i(k)$ for each seller i , we conclude $p(k+1) \cdot (k+1) = \min_{i \neq w(k+1)} v_i^S(k+1) \geq \min_{i \neq w(k)} v_i^S(k) = p(k) \cdot k$. \square

Theorem 5.6. *\mathcal{M}_{2nd} is buyer-truthful.*

Proof. The auction result of \mathcal{M}_{2nd} for buyer i is either receiving one unit with certain payment or receiving nothing with no payment. If i received one unit, then $v_i^B(1) \geq p(k^*)$ and the payment of i is $p(k^*)$ which is independent of $v_i^B(1)$. Otherwise, we know that $v_i^B(1) < p(k^*)$ and the payment is zero for i . Therefore, the first property of Lemma 5.4 is satisfied for all buyers.

In order to prove truthfulness, we need to show that the utility of each buyer is maximised, i.e. the payment is minimised, by \mathcal{M}_{2nd} . For all buyers who received a unit, the payment $p(k^*)$ is the same for all of them. If any of the winning buyers with valuation $v_i^B(1)$ reported $\hat{v}_i^B(1) < p(k^*) \leq v_i^B(1)$, this buyer will not win. Moreover, from Lemma 5.5, we know that $p(k^*)$ is minimal as k^* is maximal. Therefore, $p(k^*)$ is the minimum valuation for buyers to win in \mathcal{M}_{2nd} . Thus, the payment $p(k^*)$ for all winning buyers is minimised. This also holds for losing buyers. \square

Theorem 5.7. *\mathcal{M}_{2nd} is not seller-truthful.*

Proof. The auction result of \mathcal{M}_{2nd} for seller i is either selling k units with payment $p(k)$ for some $k > 0$ or selling nothing with no payment. For each $k > 0$, if seller

i successfully sells k units, then the payment $p(k) \cdot k$ received by i is the second lowest valuation of sellers for selling k units together and is independent of i 's type. If seller i loses, the payment is zero for i . Therefore, the first property of Lemma 5.4 is also satisfied for all sellers.

The reason why \mathcal{M}_{2nd} is not truthful for sellers is that the utilities of sellers might not be maximised. For instance, assume that k_1 and $k_1 - 1$ satisfy the condition $v_k^B(1) \geq p(k)$, and $w(k_1) = w(k_1 - 1) = i$. If $p(k_1) \cdot k_1 - v_i^S(k_1) < p(k_1 - 1) \cdot (k_1 - 1) - v_i^S(k_1 - 1)$, then i would prefer selling $k_1 - 1$ units other than k_1 units. Therefore, if i sells k_1 units with payment $p(k_1) \cdot k_1$, she is incentivised to manipulate the auction in order to sell only $k_1 - 1$ units with more utility. The manipulation will be successful if the third lowest seller valuation for selling k_1 units, say $v_j^S(k_1)$, satisfies $\frac{v_j^S(k_1)}{k_1} > v_{k_1}^B(1)$ (by simply misreporting $\hat{v}_i^S(k_1) \geq v_{k_1}^B(1)$). \square

5.4 A BB, IR and Seller-truthful MDA

In the last section, we showed that a simple second price MDA is not truthful, because sellers' utilities are not maximised. However, in this section, we find that if we simply update \mathcal{M}_{2nd} such that sellers' utilities are maximised, then buyers will sacrifice. The main update is that the determination of the trading size considers the winning seller's utility.

Second Price plus Seller Utility Maximisation MDA \mathcal{M}_{2nd}^+

Given type profile report $v = (v^B, v^S)$, assume that $v_1^B(1) \geq v_2^B(1) \geq \dots \geq v_m^B(1)$.

1. Let $w(k) = \min \operatorname{argmin}_i v_i^S(k)$ and $p(k) = \min_{i \neq w(k)} \frac{v_i^S(k)}{k}$ or ∞ if there is only one seller.

2. Let $k^* = \max\{k | v_k^B(1) \geq p(k)\}$, and $i^* = w(k^*)$.
3. Let $K = \{k | v_k^B(1) \geq p(k)\}$, and K^* is the least set such that $i^* \in K^*$ and $K^* \supseteq \{k | k = \max(K \setminus K^*) \wedge w(k) = i^* \wedge v_{\min K^*}^B(1) < \frac{v_{3rd}^S(\min K^*)}{\min K^*}\}$, where $v_{3rd}^S(k)$ is the third lowest valuation of sellers for selling k units and it is ∞ if there are less than three sellers.
4. Let $k_+^* = \max \operatorname{argmax}_{k \in K^*} (p(k) \cdot k - v_{i^*}^S(k))$.
5. The first k_+^* buyers, i.e. buyers of valuation $v_1^B, v_2^B, \dots, v_{k_+^*}^B$, receive one unit of the commodity each and each of them pays $p(k_+^*)$.
6. Seller i^* sells k_+^* units of the commodity and receives payment $p(k_+^*) \cdot k_+^*$.
7. The rest of the traders lose without payment.

k^* and the winning seller i^* of \mathcal{M}_{2nd}^+ is the same as that in \mathcal{M}_{2nd} . Set K contains all possible numbers of units that can be exchanged without sacrificing budget balance. Set K^* contains all k points that seller i^* can manipulate and force the auctioneer to choose some $k^* \in K^*$ if \mathcal{M}_{2nd} is used. The reason is that, for all $k \in K^*$ except the minimum ($\min K^*$), seller i^* is the only winner, i.e. without seller i^* , there is no other seller who can win at those trading sizes. Therefore, \mathcal{M}_{2nd}^+ chooses $k_+^* \in K^*$, as the final trading size, such that seller i^* 's utility is maximised among all $k \in K^*$. It is evident that \mathcal{M}_{2nd}^+ is also budget-balanced and individually rational.

Theorem 5.8. \mathcal{M}_{2nd}^+ is seller-truthful but not buyer-truthful.

Proof. Regarding truthfulness of sellers, firstly, their payments are independent of their valuations. Secondly, their utilities are maximised, i.e. they cannot misreport their valuations to get higher utilities. For winning seller i^* , K^* contains all winning k points where i^* is the winner and she can manipulate to get a winning

point giving her the highest utility. However, seller i^* cannot misreport to win at other trading sizes not included in K^* . This is because another seller will win at either $\min K^*$ or $\max(K \setminus K^*)$ if seller i^* chooses to not win at any point in K^* . Since \mathcal{M}_{2nd}^+ selects the winning point $k_+^* \in K^*$ that gives i^* the highest utility she could possibly get with misreporting, there is no reason for i^* to misreport. For a losing seller i , if i misreported and won at k^* , then i has to misreport $\hat{v}_i^S(k^*) \leq v_{i^*}^S(k^*) \leq v_i^S(k^*)$ and the K^* for i will be $\{i^*\}$. Therefore, i will get non-positive utility, $v_{i^*}^S(k^*) - v_i^S(k^*)$, in order to win at point k^* . If i misreported and won at a point $k' > k^*$, then i has to misreport $\hat{v}_i^S(k') \leq v_{i'}^B(1) \cdot k' \leq v_i^S(k')$ and the new unit price $\hat{p}(k')$ must satisfy that $\frac{\hat{v}_i^S(k')}{k'} \leq \hat{p}(k') \leq v_{i'}^B(1)$. Thus the utility for losing seller i to win at point k' will be $\hat{p}(k') \cdot k' - v_i^S(k') \leq 0$. Therefore, truthfulness also holds for losing sellers.

It is evident that \mathcal{M}_{2nd}^+ is not truthful for buyers because their payments $p(k_+^*) \geq p(k^*)$ (Lemma 5.5). That is, buyers of valuation $v_1^B, v_2^B, \dots, v_{k^*}^B$ could misreport their valuations to prevent seller i^* winning at any point $k_+^* < k^*$, which might give them higher utilities. \square

Proposition 5.9. *The utility loss of winning buyer i in \mathcal{M}_{2nd}^+ , compared with the utility i can achieve in \mathcal{M}_{2nd} , is not more than $\frac{k^* - k_+^*}{k_+^*}$ of the payment i can get when i participates in \mathcal{M}_{2nd} .*

Proof. According to Lemma 5.5, we get $p(k^*) \cdot k^* \geq p(k_+^*) \cdot (k_+^*)$. Therefore, for a winning buyer i of type v_i in \mathcal{M}_{2nd}^+ , i 's utility $u_{\mathcal{M}_{2nd}^+} = v_i(1) - p(k_+^*)$, while the utility i will get in \mathcal{M}_{2nd} is $u_{\mathcal{M}_{2nd}} = v_i(1) - p(k^*)$. So we get $u_{\mathcal{M}_{2nd}} - u_{\mathcal{M}_{2nd}^+} = p(k_+^*) - p(k^*) \leq \frac{k^* - k_+^*}{k_+^*} p(k^*)$. \square

5.5 Existence of (W)BB, IR and Truthful MDAs

Following the results in previous sections, we will demonstrate in this section that there are multi-unit double auctions that are (weakly) budget-balanced, individually rational and truthful. However, we will also prove that there does not exist a (weakly) budget-balanced, individually rational and truthful MDA, where both the trading size and the payment are neither seller-independent nor buyer-independent.

Proposition 5.10. *There exists a (weakly) budget-balanced, individually rational, and truthful multi-unit double auction.*

Proof. The fixed pricing MDA described in Auction 1 is BB, IR and truthful. Given a predetermined transaction price p , \mathcal{M}_{fixed} first calculates the total number k_1 of buyers whose valuations are at least p , then calculates the maximal number k^* of units that a seller can sell, with non-negative utility, under unit price p , given that $k^* \leq k_1$. After it calculates all the winning candidates of both sides, candidates from the same side win with the same probability. It is evident that this auction is budget-balanced and individually rational.

Regarding truthfulness, firstly, payment p does not depend on any trader. Secondly, all buyers whose valuation for one unit is at least p will win with the same probability with payment p , so their utilities are maximised if their winning probability $\frac{k^*}{k_1}$ is maximised. Buyer i of $v_i^B(1) \geq p$ will not report $\hat{v}_i^B(1) < p$ as i 's winning probability will be reduced. Also buyer i of $v_i^B(1) < p$ will not report $\hat{v}_i^B(1) \geq p$ because he will get a negative expected utility. Therefore, k_1 is fixed for a given type profile report and no buyer is incentivised to change it. Moreover, k^* is maximised. Thus, $\frac{k^*}{k_1}$ is maximised and buyers' utilities are maximised. A similar analysis applies to sellers. \square

Auction 1 (Fixed Pricing MDA \mathcal{M}_{fixed}). Given predetermined transaction price p and type profile report $v = (v^B, v^S)$,

1. let $k_1 = |\{i | v_i^B(1) \geq p\}|$,
2. let $k^* = \max\{k | k \leq k_1 \wedge \frac{v_i^S(k)}{k} \leq p \text{ for some } i\}$, and $k_2 = |\{i | \frac{v_i^S(k^*)}{k^*} \leq p\}|$,
3. randomly select k^* winning buyers from $\{i | v_i^B(1) \geq p\}$, i.e. each buyer $i \in \{i | v_i^B(1) \geq p\}$ wins with probability $\frac{k^*}{k_1}$,
4. randomly choose one winning seller from $\{i | \frac{v_i^S(k^*)}{k^*} \leq p\}$, i.e. each seller $i \in \{i | \frac{v_i^S(k^*)}{k^*} \leq p\}$ wins with probability $\frac{1}{k_2}$,
5. each winning buyer receives one unit of the commodity and pays p , the winning seller sells k^* units and receives payment $p * k^*$, and the remaining traders lose with no payment.

Note that \mathcal{M}_{fixed} is non-deterministic and the payment p does not depend on any trader. It is not hard to check that similar auctions with two fixed prices p_s, p_b such that $p_s \leq p_b$ and p_s is the unit price for winning sellers and p_b for winning buyers is (W)BB, IR and truthful. Other than fixed pricing MDAs, there are (W)BB, IR and truthful MDAs where payments are not predetermined. For instance, a simple variant of \mathcal{M}_{fixed} described in Auction 2 is one such mechanism and it is clear that \mathcal{M}_{single} is BB, IR and truthful. However, there is no MDA that is (W)BB, IR and truthful, given that both the trading size and the payment are neither seller-independent nor buyer-independent. We say a parameter of an MDA is seller-independent (buyer-independent) if the value of the parameter does not depend on sellers' (buyers') type reports.

Definition 5.11. Given MDA \mathcal{M} , a parameter d of \mathcal{M} , and type profile $v = (v^B, v^S)$, we say d is **trader-independent** if the value of d , denoted by $d_{\mathcal{M}}(\cdot)$, satisfies $d_{\mathcal{M}}(\hat{v}) = d_{\mathcal{M}}(\bar{v})$ for all $\hat{v}, \bar{v} \in R(v)$. We say d is **seller-independent** if $d_{\mathcal{M}}((\hat{v}^B, \hat{v}^S)) = d_{\mathcal{M}}((\hat{v}^B, \bar{v}^S))$ for all $\hat{v}^B \in R(v^B)$, all $\hat{v}^S, \bar{v}^S \in R(v^S)$. We say d is **buyer-independent** if $d_{\mathcal{M}}((\hat{v}^B, \hat{v}^S)) = d_{\mathcal{M}}((\bar{v}^B, \hat{v}^S))$ for all $\hat{v}^B, \bar{v}^B \in R(v^B)$, all $\hat{v}^S \in R(v^S)$.

A parameter of an MDA is trader-independent if and only if it is seller-independent and buyer-independent. For instance, p of \mathcal{M}_{fixed} is trader-independent, and p of \mathcal{M}_{single} is seller-independent.

Auction 2 (One-sided Pricing MDA \mathcal{M}_{single}). Given the type profile report $v = (v^B, v^S)$,

1. let p be the $\lceil \frac{m}{2} \rceil$ -th highest of $v_i^B(1)$ s, where m is the total number of buyers,
2. let $k_1 = |\{i | v_i^B(1) > p\}|$,
3. let $k^* = \max\{k | k \leq k_1 \wedge \frac{v_i^S(k)}{k} \leq p \text{ for some } i\}$, and $k_2 = |\{i | \frac{v_i^S(k^*)}{k^*} \leq p\}|$,
4. randomly select k^* winning buyers from $\{i | v_i^B(1) > p\}$, i.e. each buyer $i \in \{i | v_i^B(1) > p\}$ wins with probability $\frac{k^*}{k_1}$,
5. randomly choose one winning seller from $\{i | \frac{v_i^S(k^*)}{k^*} \leq p\}$, i.e. each seller $i \in \{i | \frac{v_i^S(k^*)}{k^*} \leq p\}$ wins with probability $\frac{1}{k_2}$,
6. each winning buyer receives one unit of the commodity and pays p , the winning seller sells k^* units and receives payment $p * k^*$, and all the rest of the traders lose with no payment.

Theorem 5.12. *There is no (weakly) budget-balanced, individually rational and truthful multi-unit double auction, where both the trading size and the payment are neither seller-independent nor buyer-independent.*

Before we give the proof of Theorem 5.12, we first prove some lemmas that are going to be used for the proof. Lemma 5.13 says that an IR and truthful MDA cannot have price discrimination. An MDA has *price discrimination* if buyers (sellers) pay (receive) different payments for identical goods or services. For instance, when two buyers pay different prices for receiving one unit of the same commodity in a deterministic MDA, this is price discrimination.

Lemma 5.13. *An individually rational multi-unit double auction with price discrimination is not truthful.*

Proof. Because of individual rationality, the expected payments for all winning buyers (sellers) must be not over (under) their valuations.¹ If the expected payments are not the same between winning buyers/sellers, then a winning buyer (seller) with high (low) expected payment will have a chance to manipulate the auction in order to get a low (high) expected payment by, for example, reporting the same valuation as that of a winning buyer (seller) receiving relatively a lower (higher) expected payment. \square

From Lemma 5.13, we conclude that an individually rational and truthful MDA must give the same (expected) payment for all winning buyers/sellers, and give no payment for all losing traders.

Lemma 5.14. *All winning sellers in a truthful multi-unit double auction sell the same expected number of units.*

Proof. According to Lemma 5.13, we know that all winning sellers receive the same expected payment for selling each unit. So their utilities will be higher if they sell more units. If the expected number of units to be sold is not the same among winning sellers, then a seller selling less units is incentivised to manipulate the auction in order to sell more units by simply misreporting his valuation as the seller selling relatively more units. \square

Proof of Theorem 5.12. We first assume that there is such MDA \mathcal{M} , and then we end up with a contradiction.

Let p_s and p_b be the payment (unit price) for winning sellers and winning buyers respectively. According to Lemma 5.14, without loss of generality, we assume that \mathcal{M} selects at most one winning seller. Assume the trading size is k . Let v_{min}^B be the

¹We consider expected payment to check price discrimination, because if an MDA is non-deterministic and it can assign different payments to winning buyers/sellers. However, if a non-deterministic MDA is individually rational and truthful, then the expected payment will be the same for all winning buyers/sellers and the prices should be randomly chosen from some range independent of winning traders' valuations. A non-deterministic MDA is not considered price discrimination if the expected payment is the same for all winning/losing buyers/sellers.

minimum valuation (for one unit) of all winning buyers, and v_{max}^B be the maximum valuation of all losing buyers ($v_{max}^B = 0$ if there is no losing buyer). Let v_{win}^S be the valuation of the winning seller for selling k units, and v_{min}^S be the minimum valuation of all losing sellers for selling k units ($v_{min}^S = \infty$ if there is no losing seller). Because of individual rationality, we have $\frac{v_{win}^S}{k} \leq p_s \leq p_b \leq v_{min}^B$. Since \mathcal{M} is truthful, we further get $p_s \leq \frac{v_{min}^S}{k}$ and $p_b \geq v_{max}^B$ and p_s and p_b should not depend on any winning trader. Therefore, if \mathcal{M} chooses any k satisfying any of the following four conditions, there will be proper payments $p_s \leq p_b$ only depending on v_{max}^B and v_{min}^S .

1. $\frac{v_{min}^S}{k} \leq v_{max}^B$,
2. $\frac{v_{min}^S}{k} > v_{max}^B$, $v_{min}^B \geq \frac{v_{min}^S}{k}$, and $v_{max}^B \geq \frac{v_{win}^S}{k}$,
3. $\frac{v_{min}^S}{k} > v_{max}^B$, $v_{min}^B \geq \frac{v_{min}^S}{k}$, and $v_{max}^B < \frac{v_{win}^S}{k}$,
4. $\frac{v_{min}^S}{k} > v_{max}^B$, $v_{min}^B < \frac{v_{min}^S}{k}$, and $v_{max}^B \geq \frac{v_{win}^S}{k}$.

For condition (1), $p_b, p_s \in [\frac{v_{min}^S}{k}, v_{max}^B]$ s.t. $p_s \leq p_b$. For condition (2), $p_b, p_s \in [v_{max}^B, \frac{v_{min}^S}{k}]$ s.t. $p_s \leq p_b$. For condition (3), $p_b = p_s = \frac{v_{min}^S}{k}$, and $p_b = p_s = v_{max}^B$ for condition (4).

In other words, \mathcal{M} chooses any k satisfying any of the above four conditions can also get payments independent of winning traders and satisfying (weakly) budget balance. Besides these four conditions, we cannot choose any k under other conditions where we can still get (weakly) budget-balanced and winning trader independent payments, given that both k and p_s, p_b are neither seller-independent nor buyer-independent.

Therefore, in order to satisfy truthfulness, M has to choose a k such that all traders' utilities are maximised. For winning buyers, they would prefer a bigger k as their payment will be lower compared to the payment with a lower k , i.e. their utilities are maximised when k is maximised. However, the winning seller might

prefer a lower value of k as her utility is not necessarily maximised with maximum k (see the proof of Theorem 5.7 for example). Thus, we may not always be able to choose a k maximising both buyers' and sellers' utilities. This contradicts the truthfulness of \mathcal{M} , i.e. buyers may be incentivised to disable the above four conditions for lower values of k , while sellers may be motivated to disable that for higher values of k . \square

5.5.1 Competitive MDAs

Corollary 5.15. *There is no (weakly) budget-balanced, individually rational, truthful multi-unit double auction that is also competitive.*

Proof. From Theorem 5.12, we know that there is no (W)BB, IR, truthful, and competitive multi-unit double auction, if both the trading size and the payment are neither seller-independent nor buyer-independent. In the following, we will prove that if the trading size or the payment of an MDA is either seller-independent or buyer-independent, the MDA will not be competitive.

If the trading size of MDA \mathcal{M} is seller-independent, say the expected trading size is k_e , then k_e must be also buyer-independent, otherwise we can always find an example that violates budget balance, individual rationality and truthfulness. For instance, each seller's unit valuation for selling any number of units is larger than the highest valuation of sellers, in which the trading size should be zero if BB, IR and truthfulness are satisfied. Therefore, given $k_e > 0$ is trader-independent, for any type profile report v with optimal trading size $k_{opt}(v)$, the competitive ratio $c = \frac{k_{opt}(v)}{k_e}$. It is clear that c is not bounded as $k_{opt}(v)$ can be any value approaching to infinite.

If the payment of MDA \mathcal{M} is seller-independent, then for any payment determined without considering sellers, there exists a case where all sellers' unit valuation for selling any number of units are higher than the payment, which means that the

trading size will be zero if \mathcal{M} is (weakly) budget-balanced, individually rational, truthful. Therefore, \mathcal{M} cannot be competitive under this condition. This result also holds when the payment is buyer-independent. \square

5.6 Summary

This chapter studied a multi-unit double auction, where each seller has an unlimited supply, for exchanging one kind of commodity. Unlike previous studies of multi-unit double auction, the chapter introduced group buying in the model. More specifically, sellers' average unit valuation is decreasing (non-increasing) as the number of units sold together increases; that is, more buyers buying the commodity together as a group from a seller will receive a higher discount.

We found that, under this model, even without considering other criteria, budget-balanced, individually rational and truthful mechanisms are hard to achieve. We showed that in Theorem 5.12 there is no budget-balanced, individually rational and truthful multi-unit double auction, given that both the trading size and the payment of the auction are neither seller-independent nor buyer-independent (referred to independence constraint in the following). Under the independence constraint, there do exist mechanisms that are budget-balanced, individually rational but one-sided truthful; that is, truthful for either buyers or sellers (see Sections 5.3 and 5.4). Without considering the independence constraint, in Section 5.5, we did find auctions that satisfy all the three criteria. Moreover, if we consider trading size (i.e. the number of units exchanged) at the same time, we demonstrated in Corollary 5.15 that there is no budget-balanced, individually rational and truthful mechanism that can also guarantee trading size.

The results in this chapter are based on the assumption that each buyer requires only one unit. As we mentioned, the results are applicable to the general case where each buyer i requires $c_i > 0$ units. For the extension, we just need to

update $v_i^B(1)$ into $\frac{v_i^B(c_i)}{c_i}$ in the results, and count the number of units for a buyer group based on buyers' c_i s other than the number of buyers in the group. For non-deterministic MDAs, e.g. \mathcal{M}_{fixed} and \mathcal{M}_{single} , the winning probability of a buyer will be based on his c_i , e.g. the winning probability of buyer i in step 3 of \mathcal{M}_{fixed} will be $\frac{k^* \cdot c_i}{k_1}$. As c_i is not part of a buyer's private information, this extension will not affect any of the properties that hold in the single-unit demand case.

As closely related work, Huang *et al.* [2002] proposed weakly budget-balanced, individually rational and truthful multi-unit double auctions, under the model where each seller (buyer) supplies (demands) a publicly known number of units, their valuation for each unit is not changing and their requirements can be partially satisfied. Chu [2009] studied a multi-unit double auction model where there are multiple commodities, each seller supplies multi-units of one commodity and each buyer requires a bundle of different commodities. They proposed a method that intentionally creates additional competition in order to get budget-balanced, individually rational and truthful mechanisms. Wurman *et al.* [1998] also considered one-sided truthful double auctions for optimising social welfare. Goldberg *et al.* [2002] studied one-sided auctions where the seller has an unlimited supply without giving any valuation or reserve price for the commodity, and their goal is to design truthful mechanisms that guarantee the seller's revenue. For group buying, Edelman *et al.* [2011] considered the advertising effect of discount offers by modelling the procedure with two periods, so traders can come back in the future after getting discounted offers. Arabshahi [2011] provided a very detailed analysis of the *Groupon* business model and Byers *et al.* [2011] showed some primary post-analysis of *Groupon*. A very earlier study of online group buying is provided by Anand and Aron [2003].

There are many questions for considering group buying in multi-unit double auction worth further investigation. Especially, if sellers have limited supply, how do we calculate their utilities, as they should have valuation for the unsold units and

the valuation for the unsold units is not the same before and after the auction, raising the further question of how to optimise social welfare and guarantee other properties in this case. For instance, a seller supplies two units with unit prices $p_1 > p_2$ for selling one and two units respectively. If we reach a situation where one unit is left for the seller, we might consider that the seller has a valuation of p_1 for this unsold unit.

Chapter 6

Online Double Auction

Chapter 4 proposed a computational-efficient mechanism for computing the optimal offline solution of a decision-independent dynamic environment in which each trader is active only one period of time and the trader's valuation does not change during that time. This chapter looks at the corresponding online double auction design problem.

This chapter shows that no deterministic online double auction exists that is truthful and competitive for maximising social welfare in an adversarial model. Further, it shows that, when the sellers are relatively static and the demand is not more than the supply, a simple, deterministic and truthful online mechanism is actually 2-competitive; that is, it guarantees that the social welfare of its allocation is at least half of the optimal social welfare. Moreover, if the demand can be predicted exactly, it is demonstrated that an online double auction in this environment can be reduced to an online one-sided auction, and the truthfulness and competitiveness of the reduced online double auction follow those of the online one-sided auction. Notably, in the second environment, a truthful online double auction that is almost 1-competitive is achievable, when buyers arrive randomly.

6.1 Introduction

Double auctions have been well studied in simple static settings, where traders are known before the auctioneer makes any decision [McAfee, 1992, Myerson and Satterthwaite, 1983, Wurman et al., 1998]. However, in most modern double auction markets, traders are arriving and departing at different times. We call these markets *online double auctions*. The main challenge for the auctioneer in an online double auction is to make decisions without knowing the traders/orders that have not yet arrived. The decisions involve an online bipartite matching (i.e. allocation) between sellers and buyers and a payment calculation.

Following the previous work in online auction design [Blum et al., 2006, Bredin et al., 2007, Parkes, 2007], this chapter makes an incremental step in this field. We focus on two important criteria, *truthfulness* and *efficiency*, for online double auction design.

We will show that there is no deterministic and truthful online double auction that is also competitive for efficiency in an adversarial setting. Then we further study the environment where sellers are relatively static compared to buyers. Within this environment, two situations will be examined: 1) where the demand (the number of buyers) is not more than the supply (the number of sellers), but is not known exactly, and 2) where the demand is predictable and not necessarily not more than the supply. We show that, in the first situation, a deterministic and truthful mechanism is 2-competitive. In the second situation, we propose a framework to reduce a truthful online double auction to a truthful online one-sided auction, and demonstrate that the competitiveness of the reduced online double auction follows that of the online one-sided auction. Especially, in the second case, a truthful online double auction that is almost 1-competitive, i.e. the social welfare of the auction's allocation is nearly optimised, is achievable when buyers arrive randomly.

During last decade, there have been substantial researches on mechanism design in different dynamic environments (see [Parkes, 2007] for a survey). Most of the previous work has focused on one-sided dynamic markets where either the supply or the demand is dynamic, e.g. Ad auctions [Mehta et al., 2007]. More importantly, the auctioneer (in most cases, the seller) in one-sided dynamic markets does not provide valuations (or reserve prices) to the commodities exchanged and is not considered to strategically manipulate the auction. However, in online double auction markets, both the supply and the demand are dynamic and both buyers and sellers are playing strategically, and the auctioneer has no control over either of them. To tackle the complexity of online double auction design, we often utilise certain accessible prior knowledge of the dynamics to get desirable online auctions [Blum et al., 2006, Bredin et al., 2007].

As closely related work, given the assumption that the valuations of traders are in a range $[p_{min}, p_{max}]$, Blum *et al.* [2006] proposed a r -competitive truthful online double auction in an adversarial setting for maximising social welfare, where r is the fixed point of $r = \frac{1}{2} \ln \frac{p_{max} - p_{min}}{(r-1)p_{min}}$. Besides that, they also considered many other criteria. Moreover, assuming that traders' available/active time period in the auction is no more than some constant K , Bredin *et al.* [2007] designed a framework to construct truthful online double auctions from truthful static double auctions, and demonstrated the performance (for maximising social welfare) of the auctions given by the framework in probabilistic settings through experiments. Nevertheless, the competitive ratio of the auctions in [Blum et al., 2006] is restricted by the valuation range $[p_{min}, p_{max}]$ and therefore can be arbitrarily large, and the truthfulness of the auctions in [Bredin et al., 2007] relies on the constant K and the competitiveness is based on experiments.

This chapter is organised as follows. The market model and related concepts are briefly introduced in Section 6.2. In Section 6.3, we show the impossibility result. Then a deterministic and truthful mechanism that is 2-competitive is proposed in Section 6.4 and a framework to reduce a truthful online double auction to a truthful

online one-sided auction is given in Section 6.5 for two different environments respectively. Section 6.6 summarizes this chapter.

6.2 Preliminaries and Notations

We consider an online/dynamic double auction market, in which a set B of **buyers** and a set S of **sellers** trade one commodity. Buyers and sellers are **traders**. We will refer to a seller as she and a buyer or trader as he. Let $T = B \cup S$ and assume that traders are independent and no trader can be both buyer and seller at the same time, i.e. $B \cap S = \emptyset$. Each trader supplies or demands a single unit of the commodity during a specific time period called the active time of the trader. Since each trader might have different active times, they might come and leave the market at different times, which causes the dynamics of the market. Given the dynamics of the market, the **auctioneer** (market owner) is challenged by making decisions without knowledge of traders who have not yet arrived.

Each trader $i \in T$ has a privately observed **type** $\theta_i = (v_i, a_i, d_i)$, where $v_i, a_i, d_i \in \mathbb{R}^+$, v_i is i 's valuation of a single unit of the commodity, and a_i and d_i are the starting point and the ending point of i 's **active time**, i.e. the arrival and departure time of i .

Due to the revelation principle [Myerson, 2008b], we will focus on mechanisms that require traders to directly report their types. However, traders do not necessarily report their true types but no early-arrival and no late-departure misreports are permitted, i.e. given trader i 's type $\theta_i = (v_i, a_i, d_i)$, his report $\theta'_i = (v'_i, a'_i, d'_i)$ satisfies $a'_i \leq d'_i$ and $[a'_i, d'_i] \subseteq [a_i, d_i]$. The intuition behind this constraint is that traders do not recognise the market before their arrival and they do not get utility for any trade happened after their true departure time. We say a seller's report (called **ask**) $\theta_i = (v_i, a_i, d_i)$ and a buyer's report (called **bid**) $\theta_j = (v_j, a_j, d_j)$ are **matchable** if and only if $v_i \leq v_j$ and $[a_i, d_i] \cap [a_j, d_j] \neq \emptyset$. In other words, a

transaction increasing (at least not decreasing) social welfare can occur between θ_i and θ_j .

Let $\theta = (\theta_i)_{i \in T}$ denote a complete type profile, and $\theta^A = (\theta_i)_{i \in S}$ and $\theta^B = (\theta_i)_{i \in B}$ be the complete ask and bid profile respectively. Let θ_{-i} be the type profile of all traders except for i .

Definition 6.1. An **online double auction (ODA)** $\mathcal{M} = (\pi, x)$ consists of an **allocation policy** $\pi = (\pi_i)_{i \in T}$ and a **payment policy** $x = (x_i)_{i \in T}$, where $\pi_i(\theta) \in [0, 1]$ indicates the probability that trader i trades successfully during his reported active time, and $x_i(\theta) \in \mathbb{R}$ determines the payment paid to or received by trader i during his reported active time.

An allocation π is **feasible** if $\sum_{i \in B} \pi_i(\theta) = \sum_{i \in S} \pi_i(\theta)$ for all B, S and θ . An ODA $\mathcal{M} = (\pi, x)$ is feasible if π is feasible. Feasibility guarantees that the auctioneer never takes short or long position in the commodity exchanged in the market. Only feasible ODAs will be discussed in this chapter.

Given trader i of type $\theta_i = (v_i, a_i, d_i)$, report profile θ' and ODA $\mathcal{M} = (\pi, x)$, let $v(\theta_i) = v_i$, and the expected **utility** of i is defined as

$$u(\theta_i, \theta', (\pi, x)) = \begin{cases} (v(\theta_i) - x_i(\theta')) \cdot \pi_i(\theta'), & \text{if } i \in B. \\ (x_i(\theta') - v(\theta_i)) \cdot \pi_i(\theta'), & \text{if } i \in S. \end{cases}$$

Definition 6.2. An ODA $\mathcal{M} = (\pi, x)$ is **truthful** (aka **incentive-compatible**) if $u(\theta_i, (\theta_i, \theta'_{-i}), (\pi, x)) \geq u(\theta_i, \theta', (\pi, x))$ for all i , all permitted misreports θ' of θ , all type profile θ .

Definition 6.3. An ODA $\mathcal{M} = (\pi, x)$ is **efficient** if \mathcal{M} maximises the expected **social welfare**

$$W(\pi(\theta)) = \sum_{i \in B} v(\theta_i) \cdot \pi_i(\theta) + \sum_{i \in S} v(\theta_i) \cdot (1 - \pi_i(\theta)) \quad (6.1)$$

for all type profile θ .

In other words, an ODA is efficient if it always allocates items to those traders who value them most highly. In a market with dynamic participants, it is often not possible for an online mechanism to guarantee efficient allocations without the knowledge of the dynamics, because the mechanism's decision-making is challenged by the uncertainty of future participants. Therefore, in the end, we measure an online mechanism's efficiency by competitive analysis, namely, we compare the social welfare obtained by an online mechanism with the maximal social welfare one can achieve offline. Given type profile θ , let $\mathbf{Opt}(\theta)$ be the **optimal allocation** giving the optimal/maximal social welfare. Note that $\mathbf{Opt}(\theta)$ is also constrained by feasibility. The following notion of competitiveness will be used to measure the efficiency of ODAs.

Definition 6.4. An ODA $\mathcal{M} = (\pi, x)$ is **c -competitive** if for any type profile θ , the expected social welfare of $\pi(\theta)$, $W(\pi(\theta)) \geq \frac{W(\mathbf{Opt}(\theta))}{c}$. We refer to c as the **competitive ratio** of \mathcal{M} for efficiency. We say that \mathcal{M} is **competitive** if \mathcal{M} is c -competitive for some constant c .

6.3 No Deterministic Online Double Auctions are Competitive

In this section, we will demonstrate that no deterministic ODA is competitive in an adversarial setting. We prove in the following that for any deterministic and truthful ODA $\mathcal{M} = (\pi, x)$, there exists a type profile θ such that the social welfare $W(\pi(\theta))$ is far from the optimal one $W(\mathbf{Opt}(\theta))$.

Theorem 6.5. *For any deterministic and truthful ODA $\mathcal{M} = (\pi, x)$ and any $c > 0$, there exists a type profile θ such that $W(\pi(\theta)) \leq \frac{W(\mathbf{Opt}(\theta))}{c}$.*

Proof. A deterministic ODA makes decisions at a bid's/ask's arrival time, departure time and/or predefined time points. If decisions are not only made at

an ask's departure time, then we can always find a type profile θ' such that the last arrived ask θ_{last} of θ' is matched by \mathcal{M} before θ_{last} departs. Let $\theta = (\theta', \theta_*)$ where $\theta_* = (v_*, a_*, d_*)$ is a bid and it arrives after θ_{last} is matched and before θ_{last} departs. Since \mathcal{M} 's decision does not depend on traders not yet arrived, θ_* will not be matched by $\pi(\theta)$ because there is no unmatched ask available. There exists a θ_* such that θ_* is matched by $Opt(\theta)$ (if $v(\theta_*)$ is sufficiently large) and $W(\pi(\theta)) \leq \frac{v(\theta_*)}{c} \leq \frac{W(Opt(\theta))}{c}$. Therefore, if $v(\theta_*) \rightarrow \infty$, c will also approach to ∞ .

If decisions are only made at an ask's departure time, there exists a type profile θ where the last arrived bid $\theta_* = (v_*, a_*, d_*)$ arrives after the second last ask departs and departs before the last ask departs. We also get $W(\pi(\theta)) \leq \frac{v(\theta_*)}{c} \leq \frac{W(Opt(\theta))}{c}$ if $v(\theta_*)$ is sufficiently large. Note that truthfulness is necessary to guarantee that all the reports are truthful so that the social welfare is correctly measured. \square

Given the above negative result, we can still search non-deterministic and competitive mechanisms or examine cases where the dynamics is limited by, say, certain prior knowledge of the future participants. For instance, we may know the total number of traders arriving in the future or traders' valuation satisfying certain known distribution. With certain prior knowledge of the participants, we are able to design dedicated ODAs with desirable properties, e.g. [Blum et al., 2006, Bredin et al., 2007].

In the rest of this chapter, We will examine two situations with some prior information. In both cases, we assume that sellers are **patient**, i.e. they are active before the first buyer's arrival until the arrival of the last buyer. In the first case, we further assume that the demand (i.e. the number of buyers) is no more than the supply (i.e. the number of sellers), while in the other case we assume that we can predict the demand. Although we assume that sellers are relatively static in these online double auctions, they are not the same as online one-sided auctions, even that considering reserve prices, because not only buyers but also sellers are playing strategically in double auctions.

6.4 A Deterministic & Competitive Online Double Auction

Although we have the negative result in Theorem 6.5, we propose a deterministic and truthful online double auction, called \mathcal{M}_{greedy} , which is 2-competitive, under the assumption that the demand is not more than the supply in this section.

6.4.1 Specification of \mathcal{M}_{greedy}

The allocation policy, called *Best-first (Bf) Allocation*, of the deterministic ODA \mathcal{M}_{greedy} greedily matches a newly arrived bid to the best unmatched ask until there is no unmatched ask left or all bids have arrived.

The Allocation Policy of \mathcal{M}_{greedy}

Initialization:

- Rank all asks θ^A in terms of their valuation (the smaller the valuation the lower the ranking position, breaking ties randomly).

Online Matching:

Upon arrival of bid θ_i^B :

- If the unmatched ask θ_j^A with the lowest ranking position is matchable with θ_i^B , match θ_i^B with θ_j^A , i.e. both θ_i^B and θ_j^A trade with probability 1.
- Otherwise, θ_i^B is unmatched.

Figure 6.1a shows an example of the greedy allocation, where dots indicate asks and bids, the value beside each dot represents the valuation of the ask or bid, and the order of bids is their arrival order (from top to bottom, breaking ties randomly). There is a line between an ask and a bid if they are matched by the allocation policy. Before we describe the payment policy of \mathcal{M}_{greedy} , let us first introduce a notion of reachability used in the payment policy.

Let $((\theta_1^{A*}, \theta_1^{B*}), (\theta_2^{A*}, \theta_2^{B*}), \dots, (\theta_k^{A*}, \theta_k^{B*}))$ be the sequence of ask-bid pairs that are matched by the Best-first Allocation in bid's arrival order, e.g. $((2, 7), (3, 4), (5, 6))$ in the example shown in Figure 6.1, we say that two matched pairs $(\theta_i^{A*}, \theta_i^{B*})$ and $(\theta_j^{A*}, \theta_j^{B*})$ are **reachable** from each other, if $i \leq j$ and for all $i \leq m < j$, bid θ_m^{B*} and ask θ_{m+1}^{A*} are matchable. For the example shown in Figure 6.1, $(2, 7)$ and $(3, 4)$ are reachable from each other, but $(5, 6)$ is not reachable from $(2, 7)$ and $(3, 4)$ because ask of valuation 5 and bid of valuation 4 are not matchable.

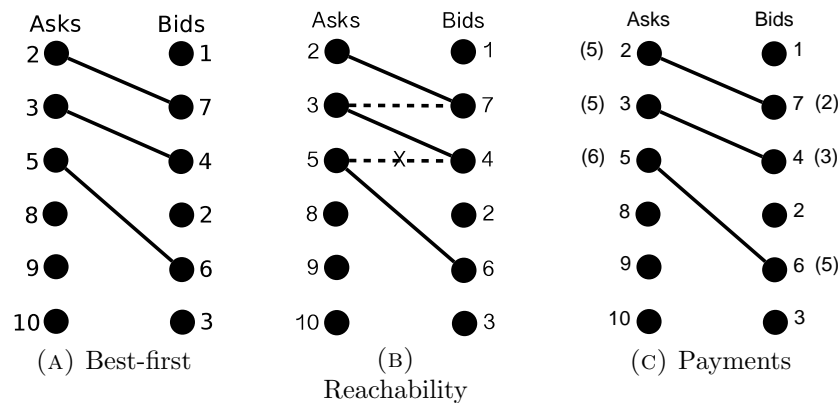


FIGURE 6.1: Running Example of \mathcal{M}_{greedy}

The payment policy is described in the following, which shows a way to calculate the VCG payment (aka *critical value* [Parkes, 2007]). Each matched buyer pays the amount equal to the valuation of the seller to whom he is matched, which is the infimum of all possible reported valuations for him to be matched in the auction, while each matched seller receives the supremum of all payments she can ask to get matched. There is no payment for unmatched traders, i.e. the mechanism is individually rational.

The Payment Policy of \mathcal{M}_{greedy}

For each matched seller i with type θ_i :

$$x_i(\theta) = \begin{cases} \min(v(\bar{\theta}_{min}^A), \max(v(\theta_{last}^B), v(\bar{\theta}_{max}^B))), & \text{if } \theta_{last}^B \text{ is reachable from } \theta_i \\ \max(v(\theta_{last}^A), v(\bar{\theta}_{max}^B)), & \text{otherwise} \end{cases}$$

where

- $\bar{\theta}_{min}^A$ is the unmatched ask that has the lowest ranking position, and $v(\bar{\theta}_{min}^A) = \infty$ if $\bar{\theta}_{min}^A$ does not exist,
- θ_{last}^A is the last matched ask,
- $\bar{\theta}_{max}^B$ is the unmatched bid that has the biggest valuation, and $v(\bar{\theta}_{max}^B) = 0$ if $\bar{\theta}_{max}^B$ does not exist,
- θ_{last}^B is the last matched bid.

For each matched buyer j with type θ_j :

$$x_j(\theta) = v(m(\theta_j))$$

where $m(\theta_j)$ is the ask matched to θ_j .

Example in Figure 6.1c shows the payments beside matched asks and bids according to the above payment rule. In this example, $v(\bar{\theta}_{min}^A)$ is 8, $v(\theta_{last}^A)$ is 5, $v(\bar{\theta}_{max}^B)$ is 3 and $v(\theta_{last}^B)$ is 6. It is easy to see that \mathcal{M}_{greedy} is running a deficit in this example. In other words, \mathcal{M}_{greedy} is not budget-balanced. Budget balance is not considered in this chapter.

6.4.2 Properties of \mathcal{M}_{greedy}

In the following, we prove that deterministic auction \mathcal{M}_{greedy} is truthful and 2-competitive.

6.4.2.1 Truthfulness

Theorem 6.6. *\mathcal{M}_{greedy} is truthful for both valuation and time arrival and departure.*

Proof. We will prove the theorem for buyers and sellers respectively.

For buyers: Since the payment for matched buyers are non-decreasing over time because of the valuation increasing of the lowest unmatched ask, the earlier the arrival time a buyer has, the higher probability to be matched and the lower payment the buyer will get. Therefore, all buyers are incentivized to arrive at their true/earliest arrival time. Since the mechanism does not use buyer's departure time for decision-making, there is no motivation for buyers to misreport their departure time.

Regarding their valuation reporting, for a matched buyer i with bid θ_i , assume $m(\theta_i) = \theta_j$, i.e. θ_i is matched to θ_j . i 's payment only depends on $v(\theta_j)$ and $v(\theta_j)$ is independent of θ_i , so the payment of i cannot be changed by $v(\theta_i)$. Moreover, increasing $v(\theta_i)$ does not change the probability for θ_i to be matched, while decreasing $v(\theta_i)$ will reduce the probability for θ_i to be matched. For an unmatched buyer i with bid θ_i , since θ_i cannot be matched to the currently best unmatched ask on the arrival of θ_i or there is no unmatched ask left, i might be able to increase his valuation to get matched, but then he has to pay more than his valuation, i.e. i gets negative utility. Thus, reporting valuation truthfully gives buyers the highest expected utility.

For sellers: All sellers are incentivized to arrive and depart truthfully as they will not be considered if they arrive after the first buyer's arrival or depart before the last buyer's arrival.

For a matched seller i with ask θ_i , we will show that i cannot report a different valuation other than her true valuation to improve her payment. Let m be the matching given by Best-first Allocation and $\theta_j = m(\theta_i)$. Assume that θ_i and θ_j is the i -th matched pair in m and $|m| = k$, i.e. the k -th matched pair of m is θ_{last}^A and θ_{last}^B . The following proof is given on the condition whether or not θ_{last}^B is reachable from θ_i .

(1) θ_{last}^B is reachable from θ_i (e.g. ask 2 in Figure 6.2a):

- If i reported θ'_i instead of θ_i such that $v(\theta'_i) > v(\theta_i)$ and $v(\theta'_i) \leq \min(v(\bar{\theta}_{min}^A), \max(v(\theta_{last}^B), v(\bar{\theta}_{max}^B)))$, the ranking position of θ'_i is $i' \geq i$ and the allocation will give a new matching m' . For all $i \leq l < i'$, the l -th matched bid of m will be matched to $(l + 1)$ -th matched ask of m in m' , θ'_i will be matched to i' -th matched bid of m in m' , and for all $1 \leq l < i$ and $i' < l \leq k$, the l -th matched pair of m is also a matched pair in m' (see Figure 6.2b and 6.2c for example). In both m and m' , the payment for i is the same because θ_{last}^B is still reachable from θ'_i , and $\bar{\theta}_{min}^A$, θ_{last}^B and $\bar{\theta}_{max}^B$ are not changed. Moreover, the probability for trader i to be matched will be the same with both θ_i and θ'_i , which is 1 here. However, if $v(\theta'_i) > \min(v(\bar{\theta}_{min}^A), \max(v(\theta_{last}^B), v(\bar{\theta}_{max}^B)))$, then θ'_i will not be matched in m' (see Figure 6.2d for example). Therefore, i cannot report a higher valuation to receive more payment.
- If i reported θ'_i instead of θ_i such that $v(\theta'_i) < v(\theta_i)$, we know that θ'_i will be matched. There will be two situations: 1) θ_{last}^B is still reachable from θ'_i , 2) θ_{last}^B is not reachable from θ'_i . In the first situation, $\bar{\theta}_{min}^A$, θ_{last}^B and $\bar{\theta}_{max}^B$ of m and m' are the same, so the payment will be the same for θ'_i and θ_i . In the second situation, we will have two sub-cases: a) $\bar{\theta}_{min}^A$ of m is θ_{last}^A of m' and

$\bar{\theta}_{max}^B$, θ_{last}^B are the same for both m and m' (see the manipulation example in Figure 6.3c and 6.3a in another way around, i.e. ask of 4.5 is misrepresented as ask of 2), b) θ_{last}^B of m is $\bar{\theta}_{max}^B$ of m' and $\bar{\theta}_{min}^A$ is the same for both m and m' (see the manipulation example in Figure 6.3e and 6.3d (or Figure 6.3f and 6.3d) in another way around). Following the proof for the condition “ θ_{last}^B is not reachable from θ_i ” in the following, we conclude that i cannot improve her utility by misreporting a lower valuation.

(2) θ_{last}^B is not reachable from θ_i (e.g. ask 2 in Figure 6.3a/6.3d):

- If i reported θ'_i instead of θ_i such that $v(\theta'_i) > v(\theta_i)$ and $v(\theta'_i) \leq \max(v(\theta_{last}^A), v(\bar{\theta}_{max}^B))$, we will get a new matching m' . If θ_{last}^B is still not reachable from θ'_i in m' , then the payment for θ'_i is the same as for θ_i (see Figure 6.3b for example). If θ_{last}^B of m is reachable from θ'_i in m' and it is also the last matched bid of m' (i.e. $v(\theta_{last}^A) > v(\bar{\theta}_{max}^B)$), then θ_{last}^A of m is $\bar{\theta}_{min}^A$ of m' and therefore the payment for θ'_i will be the same as for θ_i (e.g. Figure 6.3a and 6.3c). If $v(\theta_{last}^A) \leq v(\bar{\theta}_{max}^B)$ and θ_{last}^B of m' is reachable from θ'_i in m' , then θ_{last}^B of m' will be $\bar{\theta}_{max}^B$ of m and θ_{last}^A of m' is either θ_{last}^A of m or θ'_i (see Figure 6.3d, 6.3e and 6.3f for example). It is evident that the payment in this case is also not improved. However, if $v(\theta'_i) > \max(v(\theta_{last}^A), v(\bar{\theta}_{max}^B))$, then θ'_i will not be matched in m' . Therefore, i cannot improve her payment by reporting a higher valuation.
- If i reported θ'_i instead of θ_i such that $v(\theta'_i) < v(\theta_i)$, the ranking position of θ'_i might be lower than that of θ_i , but it will not change the probability for i to be matched, θ_{last}^A and $\bar{\theta}_{max}^B$ are still the same, and θ_{last}^B is still not reachable from θ'_i . Thus, the payment will be the same for i with both reports θ_i and θ'_i .

For an unmatched seller i with ask θ_i , we know that $v(\theta_{last}^A) \leq v(\theta_i) > v(\bar{\theta}_{max}^B)$. If i reported θ'_i such that $v(\theta'_i) < v(\theta_i)$ and θ'_i is matched in the new matching m' ,

then there will be three cases: (a) θ_{last}^A and θ_{last}^B of m are also those of m' , (b) θ_{last}^A of m is $\bar{\theta}_{min}^A$ of m' and θ_{last}^B of m is θ_{last}^B of m' , (c) $\bar{\theta}_{max}^B$ of m is θ_{last}^B of m' and either θ'_i or θ_{last}^A of m is θ_{last}^A of m' . For any of these three cases, the payment for i with report θ'_i will be less than or equal to $v(\theta_i)$, i.e. i gets non-positive utility by misreporting. \square

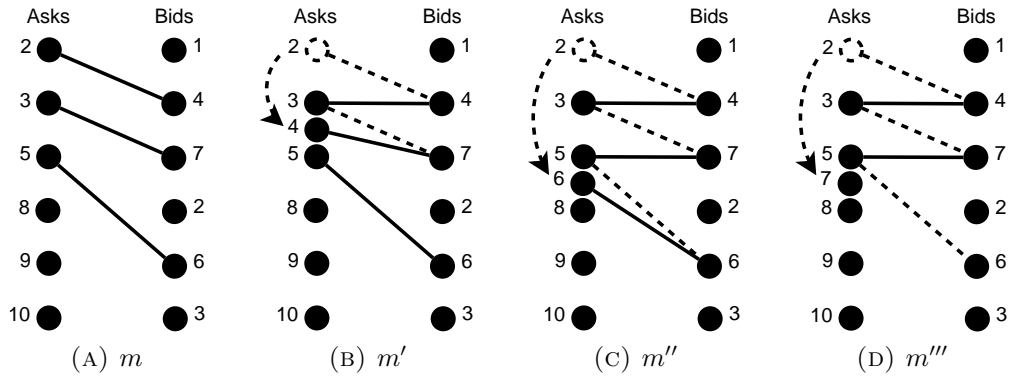


FIGURE 6.2: Seller Manipulation Examples I

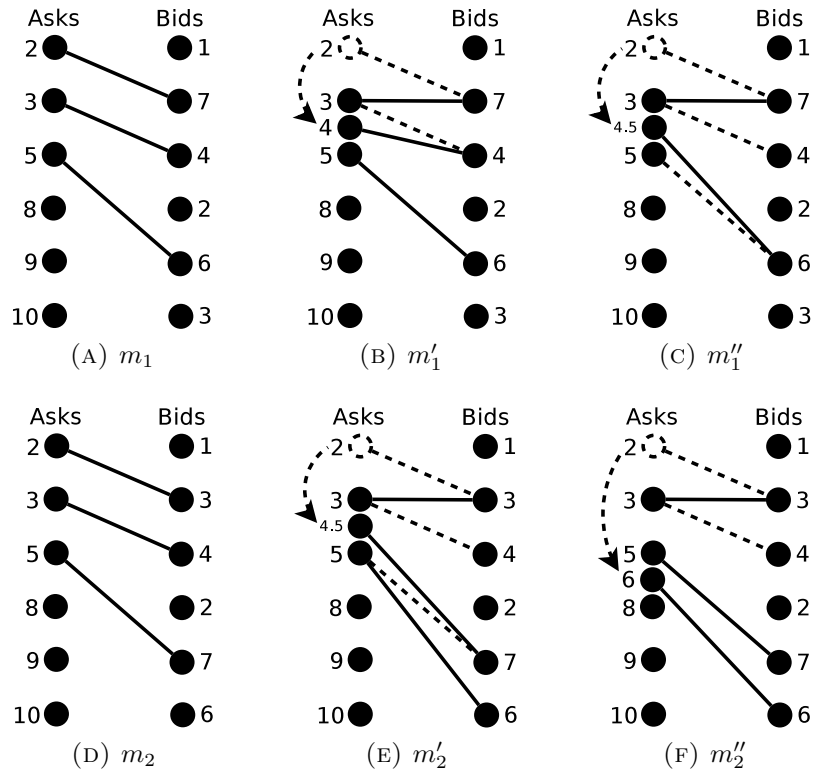


FIGURE 6.3: Seller Manipulation Examples II

6.4.2.2 Efficiency

In this section, we will apply competitive analysis, a method invented for analysing online algorithms, to check the efficiency of \mathcal{M}_{greedy} . In other words, we will determine a competitive ratio c , defined in Definition 6.4, for \mathcal{M}_{greedy} .

To that end, given a report profile θ , we need to first know what is the optimal allocation, i.e. an allocation maximising social welfare, if we are aware of all inputs/reports in advance. In this case, the optimal allocation is achieved by matching the highest bid (with respect to valuation) with the lowest ask, the second highest bid with the second lowest ask and so on, until there is no more matchable pair left. Figure 6.4 shows an example for both Best-first Allocation and optimal allocation.

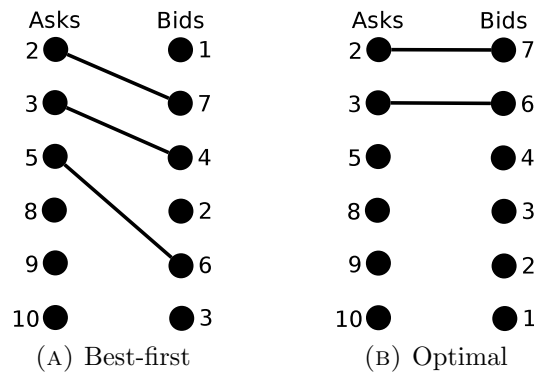


FIGURE 6.4: Best-first Allocation vs Optimal Allocation

Given the optimal allocation, we are ready to calculate the competitive ratio of Best-first Allocation. We first show that all asks that are matched by the optimal allocation are also matched by Best-first Allocation.

Lemma 6.7. *All asks that are matched by the optimal allocation are also matched by Best-first Allocation.*

Proof. We know that in the optimal allocation, any matched ask can be matched to any matched bid. Since all the matched asks of the optimal allocation will be matched first by Best-first Allocation, each of the matched asks of the optimal

allocation will be matched to either one of the matched bids or another unmatched bid of the optimal allocation by Best-first Allocation. \square




FIGURE 6.5: A Special Case of Best-first Allocation

Theorem 6.8. \mathcal{M}_{greedy} is 2-competitive.

Proof. We first show that this competitive ratio is achievable under a special case and then we prove that in any other cases the ratio is also achievable.

The special case is that all matched asks of the optimal allocation are matched to unmatched bids of the optimal allocation by Best-first Allocation, and unmatched asks of the optimal allocation are also not matched by Best-first Allocation (see Figure 6.5 for example, where the grey areas are the asks and bids matched by Best-first Allocation and double arrow lines indicate the matching relation). Let A_{Opt} and \bar{A}_{Opt} be the matched and unmatched asks respectively in the optimal allocation, and B_{Opt} and B_{Bf} be the matched bids in the optimal allocation and Best-first Allocation respectively and \bar{B}_{Opt} and \bar{B}_{Bf} be the corresponding unmatched bids. We can get that $B_{Opt} \cap B_{Bf} = \emptyset$, i.e. no bid from B_{Opt} can be matched to any ask from \bar{A}_{Opt} . We also know that $\bar{A}_{Opt} \neq \emptyset$ and $|\bar{A}_{Opt}| \geq |B_{Opt}|$ because we assumed that the demand is not more than the supply. Therefore,

$$\sum_{\theta_i \in \bar{A}_{Opt}} v(\theta_i) > \sum_{\theta_i \in B_{Opt}} v(\theta_i). \quad (6.2)$$

The social welfare of the optimal allocation is:

$$W(Opt(\theta)) = \sum_{\theta_i \in \bar{A}_{Opt}} v(\theta_i) + \sum_{\theta_i \in B_{Opt}} v(\theta_i). \quad (6.3)$$

The social welfare of Best-first Allocation is:

$$W(Bf(\theta)) = \sum_{\theta_i \in \bar{A}_{Opt}} v(\theta_i) + \sum_{\theta_i \in B_{Bf}} v(\theta_i). \quad (6.4)$$

Combining (6.2), (6.3) and (6.4), we get

$$\frac{W(Bf(\theta))}{W(Opt(\theta))} > \frac{\sum_{\theta_i \in \bar{A}_{Opt}} v(\theta_i) + \sum_{\theta_i \in B_{Bf}} v(\theta_i)}{\sum_{\theta_i \in \bar{A}_{Opt}} v(\theta_i) + \sum_{\theta_i \in \bar{A}_{Opt}} v(\theta_i)} > \frac{1}{2}.$$

So far, we have proved the theorem in a special case. In the general case, some asks of A_{Opt} might be matched to some bids of B_{Opt} , and some asks of \bar{A}_{Opt} might be matched to some bids of B_{Opt} by Best-first Allocation. Due to Lemma 6.7, we know that all asks in A_{Opt} are matched by Best-first Allocation. Let B_A and \bar{B}_A be all the bids from B_{Opt} and \bar{B}_{Opt} respectively that are matched to asks of A_{Opt} by Best-first Allocation. Let \bar{A}_B be the asks from \bar{A}_{Opt} that are matched to some bids of B_{Opt} by Best-first Allocation, and $B_{\bar{A}}$ be the corresponding bids matched to \bar{A}_{Opt} . Let $B' = B_{Opt} \setminus (B_A \cup B_{\bar{A}})$ be the asks from B_{Opt} that are not matched by Best-first Allocation (see Figure 6.6). Therefore, the social welfare of Best-first Allocation is:

$$W(Bf(\theta)) = \sum_{\theta_i \in \bar{A}_{Opt} \setminus \bar{A}_B} v(\theta_i) + \sum_{\theta_i \in B_A \cup \bar{B}_A \cup B_{\bar{A}}} v(\theta_i).$$

So, we get

$$\begin{aligned}
\frac{W(Bf(\theta))}{W(Opt(\theta))} &= \frac{\sum_{\theta_i \in \bar{A}_{Opt} \setminus \bar{A}_B} v(\theta_i) + \sum_{\theta_i \in B_A \cup \bar{B}_A \cup B_{\bar{A}}} v(\theta_i)}{\sum_{\theta_i \in \bar{A}_{Opt}} v(\theta_i) + \sum_{\theta_i \in B_{Opt}} v(\theta_i)} \\
&= \frac{\sum_{\theta_i \in \bar{A}_{Opt} \cup B_{Opt}} v(\theta_i) - \Sigma}{\sum_{\theta_i \in \bar{A}_{Opt} \cup B_{Opt}} v(\theta_i)} \\
&= 1 - \frac{\Sigma}{\sum_{\theta_i \in \bar{A}_{Opt} \cup B_{Opt}} v(\theta_i)}, \tag{6.5}
\end{aligned}$$

where $\Sigma = \sum_{\theta_i \in B' \cup \bar{A}_B} v(\theta_i) - \sum_{\theta_i \in \bar{B}_A} v(\theta_i)$.

Since the number of bids is not more than that of asks, i.e. the number of unmatched bids is not more than that of unmatched asks, we get $|\bar{A}_{Opt} \setminus \bar{A}_B| > |B'|$. We know that no ask from $\bar{A}_{Opt} \setminus \bar{A}_B$ can be matched to any bid from B' , so $\sum_{\theta_i \in \bar{A}_{Opt} \setminus \bar{A}_B} v(\theta_i) \geq \sum_{\theta_i \in B'} v(\theta_i)$, i.e. $\sum_{\theta_i \in \bar{A}_{Opt}} v(\theta_i) \geq \sum_{\theta_i \in B' \cup \bar{A}_B} v(\theta_i)$. Thus,

$$\frac{\Sigma}{\sum_{\theta_i \in \bar{A}_{Opt} \cup B_{Opt}} v(\theta_i)} \leq \frac{\Sigma}{\sum_{\theta_i \in B' \cup \bar{A}_B} v(\theta_i) + \sum_{\theta_i \in B_{Opt}} v(\theta_i)}. \tag{6.6}$$

Since

$$\sum_{\theta_i \in B_{Opt}} v(\theta_i) = \sum_{\theta_i \in B_{\bar{A}} \cup B_A \cup B'} v(\theta_i) \geq \sum_{\theta_i \in \bar{A}_B \cup B_A \cup B'} v(\theta_i) \geq \sum_{\theta_i \in \bar{A}_B \cup B'} v(\theta_i),$$

we conclude that

$$\begin{aligned}
&\frac{\Sigma}{\sum_{\theta_i \in B' \cup \bar{A}_B} v(\theta_i) + \sum_{\theta_i \in B_{Opt}} v(\theta_i)} \leq \\
&\frac{\Sigma}{\sum_{\theta_i \in B' \cup \bar{A}_B} v(\theta_i) + \sum_{\theta_i \in \bar{A}_B \cup B'} v(\theta_i)} \leq \\
&\frac{\Sigma + \sum_{\theta_i \in \bar{B}_A} v(\theta_i)}{\sum_{\theta_i \in B' \cup \bar{A}_B} v(\theta_i) + \sum_{\theta_i \in \bar{A}_B \cup B'} v(\theta_i)} = \frac{1}{2}. \tag{6.7}
\end{aligned}$$

Combining (6.5), (6.6) and (6.7), we get $\frac{W(Bf(\theta))}{W(Opt(\theta))} \geq \frac{1}{2}$. \square




FIGURE 6.6: A General Case of Best-first Allocation

6.5 Reducing Online Double Auctions to Online One-sided Auctions

In Section 6.3, we have shown that there is no deterministic mechanism that can guarantee efficiency if we do not have any information of the incoming bids. However, if we know that the demand is not more than the supply, we can use this prior knowledge to get a deterministic online mechanism that is also truthful and 2-competitive in the last section. In this section, we will study another case where we can predict how many bids will arrive. Given this prior information, we will demonstrate how to reduce an ODA to an *online one-sided auction* that aims to select the k -best bids from n bids arriving in an online fashion, e.g. secretary-problem-based online auctions [Buchbinder et al., 2010, Hajiaghayi et al., 2004, Kleinberg, 2005].

The main difference between ODAs and online one-sided auctions is that, instead of selling k items to n bidders in online one-sided auctions, we do not know how many items we should allocate to buyers in ODAs, because items are provided by sellers who have valuations on the items. In other words, it is not necessary to sell an item from a seller with high valuation to a buyer with low valuation, if maximising social welfare is part of the goal of the auction. We refer to seller's valuation as reserve price. Therefore, one way to reduce ODAs to online one-sided

auctions is integrating reserve prices in online one-sided auctions. For instance, in the situation where there is only one seller, we want to select a buyer that has a valuation at least better than the seller's. Although the buyer we selected has the highest valuation among all buyers, the buyer's valuation might be still very small compared with the valuation of the seller. What if we consider the seller in the above example as an additional buyer and let the one-sided auctions select the best among both the seller and buyers? In the rest of this section, we will show how to consider sellers as additional buyers in ODAs, and how to achieve the truthfulness and competitiveness of ODAs by using truthful and competitive online one-sided auctions.

6.5.1 The Reduction

Let n^A and n^B be the number of asks θ^A and bids θ^B respectively. Let \mathcal{A} be an online one-sided auction. We construct an ODA $\mathcal{M}_{\mathcal{A}}$ from \mathcal{A} as follows. The intuition is considering sellers as additional buyers by giving asks opportunity to compete with bids in order to gain items back for sellers, if sellers' valuations are comparatively high among the valuations of both sellers and buyers. By doing this, a seller with a comparatively high valuation will have a better chance to get her item back if maximising social welfare is an objective of \mathcal{A} . In order to treat relatively static sellers as buyers, we need to assign them an new online arrival order in terms of the arrival of buyers in the reduction.

Online Double Auction $\mathcal{M}_{\mathcal{A}}$ based on Online One-sided Auction \mathcal{A}

1. Choose a position $l_i \in [1, n^A + n^B]$ for each ask θ_i according to a discrete probability distribution function $f(x)$ that satisfies the assumptions made on the arrival order of the inputs of \mathcal{A} .

2. Run \mathcal{A} on the inputs that contain both asks θ^A and bids θ^B where each ask θ_i arrives right after the $(l_i - 1)$ -th bid or ask arrived.
3. If a bid θ_i is selected by \mathcal{A} with payment p_i and $v(\theta_i) \geq v(\theta_j)$, where θ_j is the currently unmatched ask with lowest valuation (breaking ties randomly), then θ_i is matched to θ_j with payment

$$x_i(\theta) = \max(p_i, v(\theta_j)) \quad (6.8)$$

and therefore θ_j is also matched (initially all asks are unmatched). Otherwise, θ_i is unmatched.

4. Once the matching/allocation is done, the payment for each matched ask θ_j is defined as

$$x_j(\theta) = \min(v(\theta_{last}^B), v(\bar{\theta}_{min}^A)), \quad (6.9)$$

where θ_{last}^B is the last bid in arrival order selected by \mathcal{A} , and $\bar{\theta}_{min}^A$ is the unmatched ask with lowest valuation and $v(\bar{\theta}_{min}^A) = 0$ if there is no unmatched ask.

For the probability distribution function $f(x)$ of $\mathcal{M}_{\mathcal{A}}$, we only require that $f(x)$ satisfies the assumptions made on the arrival order of the inputs of \mathcal{A} . In other words, the arrival order assigned to asks satisfies the assumptions made on the arrival order of bids. For instance, if \mathcal{A} is based on a random-ordering model, e.g. secretary-problem-based online auctions [Kleinberg, 2005], then $f(x)$ can only be a random distribution function. If \mathcal{A} is based on an adversary-ordering model, then $f(x)$ can be any distribution function. More interestingly, if \mathcal{A} has no assumption made on the arrival order of its inputs, we can utilise $f(x)$ for other purpose. In single-seller case, for example, we might push the ask to the front of the inputs to guarantee a higher expected valuation of the selected trader and therefore to further improve the efficiency of $\mathcal{M}_{\mathcal{A}}$.

Figure 6.7 shows a running example of $\mathcal{M}_{\mathcal{A}}$. $\mathcal{M}_{\mathcal{A}}$ first chooses a position for each ask, then runs \mathcal{A} on the merged input and selects the winners (indicated by ‘*’), and finally determines the final asks and bids that are matched by using the winners selected by \mathcal{A} (traders allocated an item by $\mathcal{M}_{\mathcal{A}}$ are indicated by circles). From the example in Figure 6.7, we can say that both the ask of value 2 and the bid of value 6 do not get item in the end, although they are selected by \mathcal{A} . That is, $\mathcal{M}_{\mathcal{A}}$ might improve the social welfare of the allocation given by \mathcal{A} .




FIGURE 6.7: A Running Example of $\mathcal{M}_{\mathcal{A}}$

6.5.2 Key Properties of $\mathcal{M}_{\mathcal{A}}$

In the rest of this section, we show two key properties of $\mathcal{M}_{\mathcal{A}}$ which are not dependent on the definition of $f(x)$, and then show two instances of $\mathcal{M}_{\mathcal{A}}$.

Theorem 6.9. *If \mathcal{A} is truthful for both valuation and time arrival and departure, then $\mathcal{M}_{\mathcal{A}}$ is truthful for both valuation and time arrival and departure.*

Proof. We prove for sellers and buyers respectively. We need to show that both sellers and buyers will reveal their true valuation and arrive and departure truthfully, i.e. traders are incentivized to arrive as early as they can and depart as late as possible.

For a buyer i of type θ_i that is not selected by \mathcal{A} , θ_i will also not be matched by $\mathcal{M}_{\mathcal{A}}$. If i misreported θ'_i and is selected by \mathcal{A} , then $v(\theta_i) - p_i \leq 0$, i.e. i might get negative utility in \mathcal{A} , because \mathcal{A} is truthful. Therefore, if θ'_i is matched by $\mathcal{M}_{\mathcal{A}}$, then i 's utility $v(\theta_i) - \max(p_i, v(\theta_j)) \leq v(\theta_i) - p_i \leq 0$.

For a buyer i of type θ_i that is selected by \mathcal{A} , θ_i will be either matched or unmatched by $\mathcal{M}_{\mathcal{A}}$ depending on $v(\theta_i)$ and the lowest unmatched ask θ_j when θ_i is selected. If $v(\theta_i) \geq v(\theta_j)$, θ_i is matched by $\mathcal{M}_{\mathcal{A}}$. Otherwise, θ_i is unmatched. If θ_i is matched by $\mathcal{M}_{\mathcal{A}}$, then we know that i 's utility $v(\theta_i) - \max(p_i, v(\theta_j))$ is maximised, because $v(\theta_i) - p_i$ is maximised by \mathcal{A} and $v(\theta_j)$ is independent of i and it is minimised if i arrives at his earliest arrival time. If θ_i is not matched by $\mathcal{M}_{\mathcal{A}}$, then we have $p_i \leq v(\theta_i) < v(\theta_j)$. Since $v(\theta_j)$ is independent of i and it is minimised if i arrives at his earliest arrival time, i can only be matched by $\mathcal{M}_{\mathcal{A}}$ if i misreported θ'_i such that $v(\theta'_i) \geq v(\theta_j)$, but then $v(\theta_i) - \max(p_i, v(\theta_j)) < 0$.

For buyers, we conclude from the above that they are incentivized to arrive at their earliest arrival time and report their true valuation. Moreover, $\mathcal{M}_{\mathcal{A}}$ does not use their departure time for decision-making, so the truthfulness of their departure directly follows that of \mathcal{A} .

For sellers, since we assume that all sellers are patient, i.e. sellers arriving after the arrival of the first bid or departing before the last bid's arrival will not be considered by $\mathcal{M}_{\mathcal{A}}$, all sellers are incentivized to arrive and depart truthfully. In the following, we will prove that sellers are also incentivized to reveal their true valuation in $\mathcal{M}_{\mathcal{A}}$, i.e. telling their true valuations maximises their expected utilities.

For a matched seller i with ask θ_i , her payment $x_i(\theta)$, defined by (6.9), is not dependent on θ_i . If i asks more than $x_i(\theta)$, she will not be matched, and if i asks any value in between $v(\theta_i)$ and $x_i(\theta)$, she will increase her chance to be selected by \mathcal{A} because \mathcal{A} 's goal is to choose inputs with higher valuation, which might reduce the probability for i be matched by $\mathcal{M}_{\mathcal{A}}$ as less bids will be selected by \mathcal{A} . But if i reports a valuation less than $v(\theta_i)$, he might get less payment/utility, because bids with lower valuations might be selected as asks will become less competitive if sellers misreport a lower valuation.

For an unmatched seller i of type θ_i , the payment for all matched asks, say p^A , is at most $v(\theta_i)$ because $\min(v(\theta_{last}^B), v(\bar{\theta}_{min}^A)) \leq v(\bar{\theta}_{min}^A) \leq v(\theta_i)$. In order to get matched, i needs to bid θ'_i such that $v(\theta'_i) \leq p^A$, but then the payment for i will not be more than p^A , i.e. i might get negative utility. Thus, all sellers will reveal their true valuations on their arrival. \square

Theorem 6.10. *If \mathcal{A} is c -competitive, then $\mathcal{M}_{\mathcal{A}}$ is c -competitive.*

Proof. Given report profile θ , let $A_{\mathcal{A}}$ and $B_{\mathcal{A}}$ be the sets of selected asks and bids in \mathcal{A} respectively. Since \mathcal{A} is c -competitive for maximising social welfare, we get

$$W(\mathcal{A}(\theta)) = \sum_{\theta_i \in A_{\mathcal{A}} \cup B_{\mathcal{A}}} v(\theta_i) \geq \frac{W(\text{Opt}(\theta))}{c}.$$

Based on the winners $A_{\mathcal{A}} \cup B_{\mathcal{A}}$ selected by \mathcal{A} , $\mathcal{M}_{\mathcal{A}}$ will further improve the sellers and buyers that are going to have items. More specifically, a selected ask by \mathcal{A} might be matched by $\mathcal{M}_{\mathcal{A}}$ and the item of the seller will be allocated to a seller with comparatively higher valuation (e.g. the ask of value 2 in Figure 6.7), while a selected bid by \mathcal{A} might not be matched by $\mathcal{M}_{\mathcal{A}}$ if the bid's valuation is comparatively lower (e.g. the bid of value 6 in Figure 6.7). The reason is that $A_{\mathcal{A}}$ is only used to determine at least how many sellers are going to keep their items and the $|A_{\mathcal{A}}|$ -best sellers in terms of their valuations will keep their items for sure, and that some bids of $B_{\mathcal{A}}$ might not be matched by $\mathcal{M}_{\mathcal{A}}$ if their valuations are not good enough. As such, the social welfare of the allocation given by $\mathcal{M}_{\mathcal{A}}$ is at least that of the allocation given by \mathcal{A} . That is, $W(\mathcal{M}_{\mathcal{A}}(\theta)) \geq W(\mathcal{A}(\theta)) \geq \frac{W(\text{Opt}(\theta))}{c}$. \square

Corollary 6.11. *Let k be the number of sellers, there exists a truthful ODA $\mathcal{M}_{\mathcal{A}}$ that is*

- $2\sqrt{e}$ -competitive for $k = 1$.
- $(1 + \frac{c}{\sqrt{k}})$ -competitive.

Corollary 6.11 follows the $2\sqrt{e}$ -competitive online single-item auction proposed by Buchbinder *et al.* [2010] via linear programming and the $(1 + \frac{C}{\sqrt{k}})$ -competitive online multi-item auction introduced by Kleinberg [2005], which approaches to 1-competitive as $k \rightarrow \infty$. Both these two one-sided auctions are based on secretary problems, i.e. $f(x)$ of \mathcal{M}_A will be an uniform random distribution function.

6.6 Summary

We have studied the mechanism design problem of online double auction markets where traders are participating dynamically in the markets. Due to the complexity of the dynamics caused by traders in online double auction markets, we proved that there is no deterministic and truthful online mechanism that is competitive for maximising social welfare in an adversarial model. However, this negative result does not scale to the situations where we can access certain prior information of the participants. In this chapter, we studied two environments where sellers are relatively static and some prior information of buyers is accessible. In the first environment, we assumed that the demand (i.e. the number of buyers) is not more than the supply (i.e. the number of sellers). Under this assumption, we proposed a deterministic, 2-competitive and truthful online mechanism in Section 6.4. In the second environment, given the prior information that the number of incoming buyers is predictable, we demonstrated in Section 6.5 how to reduce a truthful online double auction to a truthful online one-sided auction, and showed that the competitiveness of the reduced online double auction follows that of the online one-sided auction. By using the reduction framework proposed for the second environment, we achieved an online double auction that is almost 1-competitive. However, the mechanisms proposed in this chapter are not (weakly) budget-balanced, which is also an important factor besides truthfulness and efficiency and worth further investigation, though it is often very hard to achieve all three criteria together even in static cases [Gonen *et al.*, 2007, Myerson and Satterthwaite, 1983].

Different prior knowledge gives us different advantages for designing online mechanisms, as it reduces the dynamics in some sense. In very complex dynamic environments, without certain prior knowledge, in general it is impossible to obtain ideal mechanisms in an adversarial model. Therefore, one objective of mechanism design in such complex environments is to search desirable online mechanisms by utilising as less prior knowledge as possible. Except the situations studied in this work and other previous works, there are many other interesting cases existing in real applications worth further investigation, e.g. dynamic kidney exchange [[Ünver, 2010](#)]. Moreover, besides prior knowledge, randomisation has also played an important role in online algorithm design [[Ben-David et al., 1990](#), [Chrobak, 2008](#), [Karp et al., 1990](#)].

Chapter 7

Behaviour-based Adaptive Double Auction

I have demonstrated the online double design problem in a decision-independent dynamic environment, where traders are dynamically arriving and departing, each trader is only active in the auction for one period of time and their valuation does not change during their active time. This chapter looks at a decision-dependent dynamic environment in which the trader type is decision-dependent. Each trader in this kind of environment will be active in the market for multiple discrete periods of time and the trader's valuation is changing over time in response to the decisions of the auction. In other words, the dynamics of this kind of environment can be affected by the auction, which is different from the online model studied in the previous chapter.

Since the dynamics are decision-dependent, this chapter proposes an approach based on traders' behaviour model (or type model) to design online double auctions that are adaptive to market changes or somehow guide traders to behave in a certain way. Once the mechanism is able to guide traders' behaviour, the mechanism will be able to predict or control the dynamics so that more efficient allocations will be achievable. Due to most of the strategies adopted by traders in

the corresponding real applications being well-classified and studied in economics, the approach in this chapter analyses and utilises the behaviour model of each kind of trader, designs specific (trader-dependent) mechanisms for attracting/controlling them, and finally integrates these trader-dependent mechanisms to achieve adaptive mechanisms for any mixed environments with these traders. The evaluation will be carried out through the market simulation platform for the Trading Agent Competition Market Design Tournament (CAT Tournament). Because of the strong assumptions required for truthfulness, namely independent valuation and rationality, this chapter is not able to consider truthfulness. It focuses on the factors used to measure the success of a real double auction market, such as a stock exchange.

7.1 Introduction

An online double auction market allows multiple buyers and sellers to trade commodities at different times as they wish. The annual Trading Agent Competition (TAC) Market Design Tournament was established in 2007 to foster research in the design of double auction market mechanisms in a dynamic and competitive environment, particularly mechanisms able to adapt to changes in the environment [Cai et al., 2009, Parkes, 2007]. A CAT tournament consists of a series of games, and each game is a simulation of double auction markets including traders (buyers and sellers) and specialists (market makers). Traders are simulated and provided by the tournament organiser, while each specialist is a double auction market set up and simulated by a competitor. Traders dynamically swap between specialists to trade, while specialists compete with each other by attracting traders, executing more transactions and gaining more profit. Therefore, the CAT tournament environment simulates not only the dynamics of traders but also competition among specialists, which renders the market design particularly challenging.

Although certain winning market mechanisms under the TAC competition platform have been published [Honari et al., 2009, Niu et al., 2010, Stavrogiannis and Mitkas, 2009, Vytelingum et al., 2008], they cannot guarantee that a winning mechanism is also competitive when the environment is changed. This is also demonstrated by Robinson *et al.* through a post-tournament evaluation [Robinson et al., 2009]. They showed that most specialists are susceptible to environmental changes. *This phenomenon raises the question of how to design a competitive double auction market that is adaptive to environmental changes.*

Central to becoming a winning specialist in the CAT tournament is attracting as many good traders as possible in order to receive more good shouts, generate more efficient allocations and therefore create more profit for both traders and the market maker. This is also true for a real exchange market, as traders normally choose a market based on market liquidity (indicating the performance of the market) and the number of traders in the market (indicating whether traders can benefit from trading in the market). Moreover, in general there does not exist a universal mechanism that is attractive to all kinds of traders, which also explains why different exchange markets use different policies to target different traders in the real world. *Therefore, it is very important for a market maker to fully understand the market environment and target good customers.* A key approach to understanding the market environment is the analysis of market history so that traders' behaviour models can be recognised.

Therefore, in this chapter we propose an approach based on traders' behaviours to design competitive mechanisms that are also adaptive to environmental changes. By classifying and utilising traders' behaviour, we first design mechanisms that are competitive in environments with one kind of trader, and then integrate these trader-dependent mechanisms to obtain competitive mechanisms for any complex environment that is not known in advance. That is, the proposed mechanism is guided by the behaviour of the traders and also influences their behaviour, in order to achieve certain desired properties, say efficiency (or maximising social welfare).

This chapter is organised as follows. After a brief introduction to the CAT tournament platform in Section 7.2, we show how to classify traders based on their behaviour in Section 7.3. Section 7.4 presents a way to utilise traders' behaviour in the design process and shows an experimental example. Section 7.5 demonstrates a general extension of this approach, and Section 7.6 summarises with some suggested directions for future work.

7.2 Preliminary

This section will introduce the CAT tournament platform, called JCAT [Niu et al., 2008]. JCAT provides the ability to run CAT games. A CAT game consists of a CAT server and CAT clients including traders (buyers and sellers) and specialists (market makers). The CAT server works as a communication hub between CAT clients and records all game events and validates requests from traders and specialists. A CAT game lasts a certain number of days, say 500, and each day consists of rounds. Each trading agent is equipped with a specific bidding strategy and can only choose one specialist to trade in each day, while each specialist is a combination of policies. Traders are configured by the competition organiser, and each specialist is set by a competitor.

Each trader is configured with a private value (i.e. its valuation of the goods it will trade), a *market selection strategy* and a *bidding strategy*. The market selection strategy determines a specialist to trade in each day, and the bidding strategy specifies how to make offers. The main market selection strategies used in previous competitions are based on an n-armed bandit problem where daily profits are used as rewards to update the value function. Bidding strategies integrated in JCAT are those that have been extensively studied in the literature, namely ZIC (Zero Intelligence-Constrained [Gode and Sunder, 1993]), ZIP (Zero Intelligence Plus [Cliff and Bruten, 1997]), GD (Gjerstad Dickhaut [Gjerstad and Dickhaut, 2001]), and RE (Roth and Erev [Erev and Roth, 1998]). ZIC traders bid randomly

within constraints. ZIP traders, modified version of ZIC traders, adapt to market changes to remain competitive in the market. GD traders use market history of submissions and transactions to form their beliefs over the likelihood of a bid or ask being accepted, and use this belief to guide their bidding. Finally, RE traders are designed to mimic human game-playing behaviour in extensive form games, and their strategy relies on the profit that they were able to obtain in the most recent round of trading.

Each specialist (market owner) operates one exchange market and designs its own market rules in terms of five components/policies, namely accepting policy, clearing policy, matching policy, pricing policy and charging policy. An accepting policy determines what shouts/orders are acceptable. A clearing policy schedules clearing time during a trading day. A matching policy specifies which ask is matched with which bid for clearing. A pricing policy calculates a transaction price for each match given by matching policy. A charging policy is relatively independent from other policies and determines the charges a specialist imposes on a trading day, e.g. fees for each transaction.

7.3 Behaviour-based Trader Classification

Given an unknown environment, the key for understanding the environment is analysing traders' behaviour. Especially when the strategies adopted by traders can be clearly classified, we want to find out traders' behaviour patterns for different strategies, i.e. the relationship between traders' strategies and their behaviour. Therefore, we can distinguish traders in terms of their behaviour and apply different policies for different traders. In this section, based on JCAT, we introduce how to collect traders' behaviour-related information, define the categories of traders and finally show how to classify traders based on their behaviour.

7.3.1 Data Acquisition

In JCAT, for each trader i and each specialist s , all specialists can obtain the following trader-related historical information.

- Accepted shouts of i by s .
- Cleared/Matched shouts of i by s .

The above information is also the only information about each trader available for all specialists. The trader of a rejected shout is never revealed to any specialist, even the specialist whom the shout was submitted to. Therefore, the acceptance of a shout cannot depend on the sender's historical information. Given the above information about each trader, we need to pre-process it depending on what we need for the design process, e.g. the average clearing price for a trader in a specialist during a period of time and a trader's trading time distribution.

7.3.2 Defining Categories of Trader

Given the perfect equilibrium price p_e^* of a market¹, we classify traders into two different categories, intra-marginal and extra-marginal:

- *Intra-marginal*: A seller (buyer) i with private valuation v_i is intra-marginal if $v_i \leq p_e^*$ ($v_i \geq p_e^*$).
- *Extra-marginal*: Otherwise.

The reason for classifying traders into these two categories is that intra-marginal traders can bring profitable shouts to a market, while extra-marginal traders do not. Therefore, a competitive specialist needs to attract more intra-marginal traders. We can further classify intra-marginal traders in terms of their bidding strategies.

¹The equilibrium of a market where traders truthfully report their demand and valuations.

7.3.3 Category Recognition from Trader's Behaviour

We say a trader is attracted by a specialist if the trading time the trader spent in the market of that specialist is much greater than the time it spent in any other market. We know that a profit-seeking trader chooses a specialist that has given it the highest profit in some past period. In order to give a trader profit, a specialist has to match its shouts as many as possible with profitable clearing prices. Therefore, intra-marginal traders are more likely to be attracted. Thus, a trader's trading time distribution (i.e. stability) will be the main information to be considered in its category recognition.

7.3.3.1 Trading Time Distribution

As the main market selection strategy adopted in CAT competitions, ϵ -greedy selection determines what is the most profitable specialist for a trader and then selects this specialist with probability of $1 - \epsilon$ and the others randomly with probability ϵ . This selection strategy uses reinforcement learning method based on the profit a trader received from each specialist. ϵ is mostly set to be 0.1 in CAT competitions.

Based on the above market selection strategy, we recognise the following trading time distribution patterns. We say a trader i is more stable if the time (w.r.t. the number of days) that i spent in each market varies significantly, i.e. the standard deviation of the trading time is higher. Generally speaking, intra-marginal traders are much more stable than extra-marginal traders under the same bidding strategy, but the degree of stability varies with bidding strategies.

- *Under the same bidding strategy.* All intra-marginal traders have similar trading time distribution, in other words, intra-marginal traders with valuations far from the perfect market equilibrium are not more stable than those with valuations close to the perfect market equilibrium. Extra-marginal

traders with valuations close to the perfect equilibrium are less stable than intra-marginal traders, but they still have preferences between markets. When valuations of extra-marginal traders are far from the perfect market equilibrium, they have no strict preference for any market, i.e. the times spent in each specialist are very close to each other.

- *Degree of stability with different strategies.* Given similar valuations, GD, ZIP and ZIC traders are more stable than RE traders. One reason is that an RE trader uses the profit that it was able to obtain in the most recent trading in a market to adjust (increase) its bidding price, so it will keep increasing its bidding price in a market until finally its shouts cannot be successfully matched, which will cause the trader to move to another market.

7.3.3.2 Stability vs Intra-marginality

As we have mentioned in the above, most intra-marginal traders are very stable. However, some extra-marginal traders with valuations close to the perfect equilibrium can also be very stable if there are some specialists that have very high probability to match their shouts while others cannot do so. Therefore, a stable trader doesn't need to be intra-marginal. To find out whether or not a stable trader is intra-marginal, we need further information about their behaviour, e.g. bidding prices. If a stable seller's (buyer's) average bidding price is above (under) the equilibrium price, then it maybe not intra-marginal. In general, the selected information should be able to efficiently classify traders into the categories you defined.

7.4 Behaviour-based Policy Design

A mechanism/market of a specialist is a combination of different policies and the relationship between these policies are not completely clear, so searching a competitive combination without restriction under this setting will be computationally intractable. In general, we limit the search space for each policy to certain well-known alternatives that are normally trader-independent. Moreover, there exist many policy combinations that are competitive under the same market environment, which can be seen from the results in [Niu et al., 2010]. However, in our approach, since we have gained an understanding of traders' behaviour, we are able to further limit the search space by utilising traders' behaviour. More importantly, we want to further utilise traders' behaviour information to design trader-dependent mechanisms that attract one kind of trader, and integrate those trader-dependent mechanisms to achieve adaptive mechanisms that are attractive to all kinds of traders. In the rest of this section we will define the policies of a specialist by using traders' behaviours and propose a two-step method to search adaptive mechanisms.

7.4.1 A Search Space of Behaviour-based Policies

Combined with traders' behaviours, the following policies are adapted from the literature.

7.4.1.1 Accepting Policy

Once a specialist receives a new shout, it has to first decide whether or not to accept it. If too many extra-marginal shouts are accepted, they will not be matched and therefore the transaction rate will be very low. So why does not a specialist only accept shouts from traders that it wants to attract? Unfortunately, a specialist does not know who is the sender of a shout before the shout is accepted in CAT

competitions. Instead some other general market information can be used here, e.g. the equilibrium price of historical shouts received in a market. We will use the equilibrium price of historical shouts to set up a maximum (minimum) acceptable ask (bid) price for each day, as historical equilibrium can approximately distinguish between intra-marginal and extra-marginal shouts.

Given the current day t , the most recent M historical shouts H_t^M , the maximum acceptable ask price A_t^a and the minimum acceptable bid price A_t^b are defined as:

$$\begin{aligned} A_t^a &= E(H_t^M) + \theta^a * F_t^a \\ A_t^b &= E(H_t^M) - \theta^b * F_t^b \end{aligned}$$

where $E(H_t^M)$ is the equilibrium price of H_t , $F_t^a, F_t^b \geq 0$ are relaxations, and $\theta^a, \theta^b \in [0, 1]$ are the relaxation rates. F_t^a and F_t^b are calculated for each day, and θ^a, θ^b are dynamically updated during a day, say, updated after each round.

7.4.1.2 Matching Policy

The two commonly used matching policies are *equilibrium matching* and *maximal matching* (see Chapter 3). Equilibrium matching is used to find the equilibrium price p_e which balances the bids and the asks going to be matched so that all the bids with price $p \geq p_e$ and all the asks with price $p \leq p_e$ are matched [Friedman and Rust, 1993]. The aim of maximal matching is to maximise the number of transactions/matches by matching high intra-marginal shouts with lower extra-marginal shouts if necessary. The main difference between these two matchings is that maximal matching moves some profit from high intra-marginal traders to lower extra-marginal traders so that lower extra-marginal traders are attracted. Actually maximal matching can also be used for other proposes, e.g. stabilising some high intra-marginal traders, which can be seen in a mechanism for attracting GD traders in Section 7.4.3. But one disadvantage of maximal matching is that it will heavily reduces the profit for high intra-marginal traders, and therefore they

will leave the market so that other intra-marginal traders will be affected recursively. At the same time, since equilibrium matching always gives more profit to high intra-marginal traders, some profit seeking traders, like ZIC and RE traders, will keep increasing their profit margin so that their shouts are difficult to match. Because of the availability of each traders' behaviour information, we will adopt this information for the matching policy. The following are the two additional policies we used in this framework.

1. *Double Equilibrium Matching.* We run two matchings one after another. The first matching is an equilibrium matching based on the bidding price of shouts. The second matching rematches the matched shouts given by the first matching in terms of the average clearing price of each sender's current best market², called best clearing price. The second matching matches two shouts if the gap between their best clearing prices is very small. This is because their best clearing prices are good enough to attract them and also don't give them too much space to increase their profit margin.
2. *Behaviour-based Maximal Matching.* Maximal matching is guided by the traders' behaviours so that extra-marginal shouts are matched only if the senders are those whom we want to attract, i.e. stable traders.

7.4.1.3 Pricing Policy

The pricing policy will also play a very important role not only in attracting traders but also in stabilising traders. We use a modified *discriminatory k-pricing policy*, where k is dynamically determined for each match according to the two corresponding traders' behaviour. Let $p(x)$ indicate the bidding price of shouts x , $s(x)$ indicate the sender of shout x , $best(t)$ indicate the current best market of trader t , and $p^*(t)$ is the average clearing price for trader t in $best(t)$. Assume

²The current best market of a trader is the market where it trades most.

Algorithm 7.4.1: Modified Discriminatory k -pricing Policy

Input: a : ask, b : bid**Output:** \hat{p} : clearing price

```

1 begin
2   if  $best(s(a)) = m_i$  and  $best(s(b)) = m_i$  then  $k = 0.5$ ;
3   else if  $best(s(a)) = m_i$  (or  $best(s(b)) = m_i$ ) then
4     if  $s(b)$  (or  $s(a)$ ) is attractable then  $k = minK$  (or  $k = 1 - minK$ );
5     else  $k = 0.5$ ;
6   else
7     if  $s(a)$  is more attractive than  $s(b)$  then  $k = 1 - minK$ ;
8     else  $k = minK$ ;
9   end
10  if  $p^*(a) \leq p^*(b)$  then  $p_a = \max(p^*(a), p(a))$ ;  $p_b = \min(p^*(b), p(b))$ ;
11  else  $p_a = p(a)$ ;  $p_b = p(b)$ ;
12   $\hat{p} = p_a + k * (p_b - p_a)$ ;
13 end

```

the current specialist is m_i , Algorithm 7.4.1 gives the pseudo-code of the modified pricing policy, where $minK \in [0, 1]$ is set up for different goals. The key idea of this policy is stabilising/keeping traders a specialist has already attracted and attracting those that are not attracted yet. The attractability of a trader is dependent on the overall design goal.

7.4.1.4 Clearing Policy

There are two main clearing policies used in TAC competitions, round-based and continuous. Round-based clearing clears at the end of each round, while continuous clearing clears whenever there is a new match available. The matching policy is sensitive to clearing policy. For instance, maximal matching will be useless with continuous clearing. Moreover, traders will have chances to revise their shouts if the market does not clear for some rounds during a day. We use a modified version of round-based clearing policy in this framework. Instead of clearing in each round, we choose a fixed number of clearing time points according to the number of goods each trader has, for example, we clear 5 times a day if each trader requires to exchange 3 items. Then we distribute clearing time points into

the 10 rounds of a day by giving greater preference to the first 5 rounds. Thus, we clear more in the beginning of a day while waiting longer near the end of a day, because the number of intra-marginal traders become less and less when it is approaching the end of a day and we want to give unsatisfied traders more chances to improve their shouts.

7.4.1.5 Charging Policy

Charging is a trade-off between traders' profits and a specialist's profit. It is not closely related to the above policies, but it affects the traders' market selection. Therefore, most specialists in previous competitions do not charge in the beginning of a TAC game in order to attract traders. However, for most high intra-marginal traders, charging does not affect their profit heavily, because they already reserved a large profit margin by bidding a very low (high) price to buy (sell). This framework will only focus on profit fee, as other fees, i.e. registration fee, transaction fee and information fee, could lead to 0 profit even for a trader who has successfully traded in the market.

7.4.2 Searching Adaptive Mechanisms

We know the main challenge for stabilising/attracting traders is stabilising their bidding prices, which depends on their bidding strategies. In other words, we might not be able to find a uniform mechanism that is attractive to traders with any kind of bidding strategy. Therefore, instead of searching for competitive mechanisms in a mixed environment from the very beginning, we propose a two-step approach. We first identify trader-dependent mechanisms that are competitive in an environment with only one kind of trader. Then we combine trader-dependent mechanisms together to achieve mechanisms that are competitive in any environment.

Algorithm 7.4.2: Searching Trader-dependent Mechanism

Input: m_0 : initial mechanism, f_m : a function of mechanism to maximise, δ : the minimum improvement

Output: m^* : the local best mechanism

```

1 begin
2    $CurrBest \leftarrow m_0$ ;
3   repeat
4      $m^* \leftarrow CurrBest$ ;
5     foreach policy parameter  $r$  do
6        $m' \leftarrow$  monotonically update  $r$  in  $m^*$ ;
7       if  $f_m(m') > f_m(CurrBest)$  then  $CurrBest \leftarrow m'$ ;
8     end
9   until  $f_m(CurrBest) < f_m(m^*) + \delta$  ;
10 end

```

7.4.2.1 Trader-dependent Mechanism Design

Given the goal of a trader-dependent mechanism that we want to achieve (or a function of trader-dependent mechanism to maximise), we first set up the testing environment according to the goal and an initial mechanism as the current best mechanism, and then monotonically modify only one of the parameters in the search space to compete with the current best to find the next best one that increases the value of the goal function the most, until we cannot find any modification that has any significant improvement of the function. Note that we require the modification of each parameter to be monotonic, i.e. update/change in one direction. Algorithm 7.4.2 describes the searching process for trader-dependent mechanisms. This algorithm will return mechanisms that locally maximise the goal function. In order to get an overall optimal mechanism, we can repeat this process with different initialisations.

7.4.2.2 Adaptive Mechanisms with Trader-dependent Mechanisms

Once we obtain trader-dependent mechanisms for each kind of trader offline, we adapt them online to any market environment. The main idea is to use the classification learned in Section 7.3 to determine each trader's category and to apply

the corresponding trader-dependent mechanism. However, we might end up with two inconsistent trader-dependent mechanisms that are required to run together for some environments. In such a case, we have to either apply only one of the two mechanisms or give higher priority to one of them. In order to make such a discrimination, we need to ascertain which trader-dependent mechanism will attract more good traders, which can be done, for example, by statistical analysis of traders' behaviour.

7.4.3 Experiments

In this section, we show a trader-dependent mechanism that is attractive to intra-marginal traders with the GD bidding strategy, which is also the most attractive bidding strategy adopted by traders [Phelps et al., 2010].

GD traders use the market history of submissions and transactions to form their beliefs over the likelihood of a bid or ask being accepted, and use this belief to guide their bidding [Gjerstad and Dickhaut, 2001]. Then the bidding strategy is to submit the shout that maximises a trader's expected profit, i.e. the product of its belief function and its linear utility function.

Based on the search space given in Section 7.4.1 and our specialist agent *Jackaroo*³, we identify a trader-dependent mechanism that is very good at attracting intra-marginal GD traders. The value of each parameter of the mechanism is given in Table 7.1, where A_r and B_r are respectively the accepted asks and bids until round r in one day. We have tested this trader-dependent mechanism (JaGD) with other competitive agents available from the TAC agents repository⁴, *CUNY.CS.V1* (Cu09.1), *CUNY.CS.V2* (Cu09.2), *Mertacor* (Me09), *cestlavie* (Ce09), *jackaroo* (Ja09) from CAT 2009 final, and *PoleCat* (Po10), *Mertacor* (Me10) from CAT 2010 final. Tables 7.2 and 7.3 show the average trading time distribution of one

³*Jackaroo* has achieved 1st, 2nd and 1st in CAT Tournament 2009, 2010 and 2011, respectively.

⁴<http://www.sics.se/tac/>

Policy	Parameter	Value
Accepting	F_t^a, F_t^b	6
	θ^a	$1 - \max(0, \frac{A_r - B_r}{A_r})$
	θ^b	$1 - \max(0, \frac{B_r - A_r}{B_r})$
Matching	Behaviour-based Maximal Matching	
Pricing	$minK$	0.15
Clearing	Modified Round Based	
Charging	12% profit fee	

TABLE 7.1: GD Attractive Mechanism

	Specialists								Standard Deviation
	Cu09.1	Cu09.2	Me09	Me10	Po10	Ce09	Ja09	JaGD	
ZIC Sellers	39.60	40.40	54.20	136.27	56.83	55.07	42.87	74.77	31.95
ZIC Buyers	41.77	36.13	46.23	125.53	67.13	53.93	46.17	83.10	29.64
ZIP Sellers	15.43	16.77	50.00	179.20	62.50	50.83	<u>59.40</u>	65.87	51.04
ZIP Buyers	18.30	21.07	49.00	197.83	64.50	40.90	45.37	63.03	57.24
GD Sellers	20.73	22.46	49.29	77.80	<u>87.43</u>	62.37	37.84	142.09	40.21
GD Buyers	22.91	19.57	51.23	69.50	79.84	<u>69.66</u>	41.34	145.94	40.26
RE Sellers	53.10	47.31	53.59	89.76	69.46	67.90	55.91	62.97	13.43
RE Buyers	<u>55.19</u>	<u>51.56</u>	<u>55.61</u>	86.56	73.07	64.94	55.07	58.00	11.91

TABLE 7.2: Average Trading Time Distribution of Each Type of Trader

CAT game (500 days), where the bold value in each row shows which market the traders in this row selected most and the underlined value in each column indicates which kind of traders were attracted most by the specialist in that column. The environment is mixed with 70 GD, 70 RE, 30 ZIC, and 30 ZIP buyers and sellers respectively, with valuations uniformly distributed in $[60,160]$, i.e. the perfect market equilibrium is 110. From Table 7.2 we can see that *JaGD* attracted about 30% of GD traders' trading time (the average for each market is 12.5%). Table 7.3 further shows that most traders attracted by *JaGD* are intra-marginal GD traders, and some lower extra-marginal traders are also attracted because of the use of maximal matching. It is worth mentioning that, except GD traders, this trader-dependent mechanism is not appealing to other traders, while *Me10* is good at attracting other traders but not GD traders.

Specialists									Standard Deviation
Cu09.1	Cu09.2	Me09	Me10	Po10	Ce09	Ja09	JaGD		
<i>intra-marginal buyers (with valuations between 160 and 110)</i>									
ZIC	27.60	17.27	34.13	181.27	65.27	46.20	36.73	91.53	53.39
ZIP	18.70	20.85	42.10	256.15	58.85	27.40	40.90	35.05	79.31
GD	23.97	18.64	36.72	70.92	75.42	64.81	18.53	191.00	57.02
RE	42.53	39.88	48.84	113.66	82.91	68.22	47.84	56.13	25.12
<i>lower extra-marginal buyers (with valuations between 110 and 90)</i>									
ZIC	48.25	49.38	56.13	79.38	71.88	58.75	50.38	85.88	14.62
ZIP	16.60	23.40	46.00	89.80	71.80	60.20	34.20	158.00	45.77
GD	28.44	21.44	38.89	47.67	108.89	65.22	33.56	155.89	46.81
RE	68.86	57.86	55.14	66.00	70.29	62.71	59.14	60.00	5.43
<i>other extra-marginal buyers (with valuations between 90 and 60)</i>									
ZIC	64.71	61.43	60.86	58.86	65.71	65.00	61.57	61.86	2.39
ZIP	18.40	19.60	79.60	72.60	79.80	75.60	74.40	80.00	26.99
GD	19.40	20.24	76.56	75.32	75.76	78.24	77.00	77.48	26.36
RE	65.16	62.19	62.71	63.23	63.55	62.06	61.61	59.48	1.64

TABLE 7.3: Average Trading Time Distribution of Buyers

7.5 A Framework for Behaviour-based Mechanism Design

In this section, we summarise our behaviour-based design approach to obtain a more general adaptive mechanism design framework based on traders' behaviour. This framework consists of *data acquisition*, *behaviour-based classification of traders*, *defining behaviour-based policies*, *trader-dependent mechanism design* and *integrating trader-dependent mechanisms*.

1. *Data acquisition* collects and aggregates market information, especially trader related information, which will be the foundation of the other components. Some statistical and data mining methods can be adapted here.
2. *Behaviour-based classification of traders* distinguishes traders in terms of their behaviour. This step heavily depends on the information obtained in the first step. Some machine learning methods, e.g. decision tree learning, might be useful here.

3. *Defining behaviour-based policies* determines how to utilise behaviour in specialist policies. The main contribution of traders' behaviour in this stage is connecting the loosely coupled policies to reduce the search space.
4. *Trader-dependent mechanism design* identifies mechanisms that are competitive in environments with only one of kind of trader.
5. *Integrating trader-dependent mechanisms* combines all trader-dependent mechanisms to achieve mechanisms that are competitive under an environment containing a mixture of any kinds of traders.

7.6 Summary

We have introduced a behaviour-based adaptive mechanism design approach, based on the Trading Agent Competition Market Design platform, for a dynamic double auction environment, where each trader can be active in many discrete time periods with different valuations, depending on the environment and the decisions of the mechanism. This approach consists of behaviour-based trader classification, mechanism design for specific environments (called trader-dependent mechanism design) and integrating trader-dependent mechanisms for any complex environments that are not known in advance. To the best of my knowledge, this is the first market design framework depending on traders' behaviour (or market history) to learn the market environment and guide market decisions. By integrating traders' behaviour into market policies, we are able to constrain the search space of double auction mechanisms. More importantly, because of gaining an understanding of the market environment, the resulting mechanisms will apply differential policies for attracting different traders and therefore be more focused, more competitive and adaptive. The results have been applied in the CAT player *jackaroo* which demonstrated the advantage of this approach in the Trading Agent Competition Market Design Tournament.

Chapter 8

Conclusion and Future Work

This thesis is about designing games or mechanisms that lead to socially desirable outcomes, in environments where the participants are self-interested and hold private information (type) that is the source of the outcome decision-making of the mechanism, and hence this is about *mechanism design* in general. Mechanism design has focused on how to incentivise participants to reveal their truthful private information so that desirable outcomes can be achieved. But most of the studied environments are static or simple dynamic ones, because the information uncertainty brought by a dynamic environment makes achieving certain desirable outcomes impossible in the dynamic environment. However, there are many real dynamic environments that need better mechanisms and which are not yet well-studied. In order to address this gap, this thesis offered major contributions to the mechanism design of two types of dynamic bilateral trading environment.

8.1 Summary of the Major Contributions

This thesis studied two kinds of dynamic bilateral trading environment. One type is decision-independent, where each trader's type is independently observed and

therefore the decisions of the auction cannot change it, and the other is decision-dependent, where traders' types depend on each other and also vary in response to the decisions of the auction.

For the first environment, this thesis considered a situation where each trader is only active in the market for one period of time and during this period that trader's valuation is fixed. Under this environment, this thesis proposed a reduction framework to build online double auctions from online one-sided auctions. It was shown that the truthfulness and competitiveness of these reduced online double auctions match those of the online one-sided auctions. This thesis also proposed a dedicated corresponding optimal (offline) solution by using augmentation techniques from bipartite matching. This optimal solution is one kind of VCG mechanism as it is truthful, efficient and individually rational, but it shows a very significant computational advantage; namely, it is $O(n)$ times faster than the classical VCG mechanism. Moreover, the augmentation-based approach can be extended to similar environments with constraints that are other than the temporal one considered in this thesis.

For the second environment, this thesis considered a situation where each trader can be active in many discrete time periods with different valuations, depending on the environment and the decisions of the mechanism. To address the auction design problem in this environment, this thesis proposed a behaviour-based framework to design adaptive online double auctions that can quickly adapt to the changes/dynamics of the environment. This framework designs mechanisms that first learn the behaviour model of different kinds of traders from the environment and then use the learnt results to guide the decision-making of the mechanism in order to achieve desirable allocations.

As a very good example of the second environment, a fast growing online shopping platform, which leverages group buying, has also been studied. One key reason for the success of current group buying shopping platforms such as *Groupon* is

the advertising effect. Advertising plays the major role in the dynamics of this environment, in the sense that buyers may return to buy a product at a regular price after having tried the product at a discounted price, and may recommend the product to their friends. However, sufficient modelling of the dynamics in this environment is not an easy task, and currently no good solution is extant. Nevertheless, this thesis has demonstrated that, even without the advertising effect, there is no mechanism that is truthful, individually rational and budget-balanced, if the payment and the transaction size are not predetermined. Although there do exist simple, truthful, individually rational and budget-balanced mechanisms in this model, no such mechanism exists that guarantees the number of transactions.

8.2 Future Work

Additional dynamic bilateral trading environments exist that have not been fully examined, for example, kidney exchange [[Ünver, 2010](#)].

Regarding group buying, one further interesting direction is the modelling of the dynamics of traders, especially the advertising effects and, of course, the corresponding auction design task. For example, Edelman *et al.* [[2011](#)] is trying to model the environment with two periods but in a static manner because they assumed that all buyers share a common probability of returning in the second period after their purchases in the first period. Their focus was the profitability of the seller rather than other traditional goals of mechanism design such as truthfulness and efficiency. In a real situation, the probability of a consumer purchasing an additional product, after his/her first purchase, depends on many factors, including the consumer's valuation and the prices of the product in both periods. Moreover, if a consumer is satisfied with a product, that consumer might recommend it to friends, which adds another type of dynamics or uncertainty. To address all of these factors, traditional mechanism design techniques might not be enough.

Another direction is applying the results in real markets. Although the proposed mechanisms have shown nice properties, we have also learned a lesson from Vickrey auction that theoretically beautiful mechanisms might not be applicable in real markets [[Rothkopf, 2007](#)]. There are many reasons why they are not practical. For example, human traders have bounded rationality while we assume they are completely rational in theory; traders' valuations have complicated connections which have been simplified in theory; and some properties, such as market stability, that are very important in practice, are not emphasised in theoretical study. Therefore, to apply the results to real markets, we need to adapt them to each individual environment by relaxing some of the properties considered in theory and caring additional properties of the environment.

Appendix A

Published Work

Part of the results presented in Chapters 3, 4, 5 and 7 have been published in the following papers:

- Dengji Zhao, Dongmo Zhang, Md Khan, and Laurent Perrussel: Maximal Matching for Double Auction. In the Proceedings of the 23rd Australasian Joint Conference on Artificial Intelligence (AI'10: 516–525). (The Best Student Paper Award). *Part of the results of this paper is included in Chapter 3.*
- Dengji Zhao, Dongmo Zhang and Laurent Perrussel: Mechanism Design for Double Auctions with Temporal Constraints. In the Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI'11: 472-477). *The results in this paper are contained in Chapter 4.*
- Dengji Zhao, Dongmo Zhang and Laurent Perrussel: How to Make Specialists NOT Specialised in TAC Market Design Competition? Behaviour-based Mechanism Design. In the Proceedings of the 12th International Conference on Electronic Commerce and Web Technologies (EC-Web'11: 124-135). *The results in this paper are presented in Chapter 7.*
- Dengji Zhao, Dongmo Zhang and Laurent Perrussel: Multi-unit Double Auction under Group Buying. In the Proceedings of the 20th European Conference on Artificial Intelligence (ECAI'12). *The results in this paper are presented in Chapter 5.*

Bibliography

- Anand, K. S. and Aron, R. (2003). Group buying on the Web: A comparison of price-discovery mechanisms. *Management Science*, 49(11):1546–1562.
- Arabshahi, A. (2011). Undressinggroupon: An analysis of thegroupon business model. [Available at www.ahmadalia.com/download/Undressing-Groupon.pdf; accessed 8-May-2012].
- Babaioff, M., Blumrosen, L., and Roth, A. (2010). Auctions with online supply. In *Proceedings of the 11th ACM Conference on Electronic Commerce*, EC'10, pages 13–22.
- Babaioff, M., Immorlica, N., Kempe, D., and Kleinberg, R. (2008). Online auctions and generalized secretary problems. *SIGecom Exchanges*, 7(2):1–11.
- Bar-Yossef, Z., Hildrum, K., and Wu, F. (2002). Incentive-compatible online auctions for digital goods. In *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA'02, pages 964–970.
- Ben-David, S., Borodin, A., Karp, R., Tardos, G., and Wigderson, A. (1990). On the power of randomization in online algorithms. In *Algorithmica*, pages 379–386.
- Bergemann, D. and Välimäki, J. (2010). The dynamic pivot mechanism. *Econometrica*, 78(2):771–789.
- Blum, A. and Hartline, J. D. (2005). Near-optimal online auctions. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA'05, pages 1156–1163.
- Blum, A., Kumar, V., Rudra, A., and Wu, F. (2003). Online learning in online auctions. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA'03, pages 202–204.

- Blum, A., Sandholm, T., and Zinkevich, M. (2006). Online algorithms for market clearing. *Journal of the ACM*, 53(5):845–879.
- Borodin, A. and El-Yaniv, R. (1998). *Online computation and competitive analysis*. Cambridge University Press.
- Bredin, J., Parkes, D. C., and Duong, Q. (2007). Chain: a dynamic double auction framework for matching patient agents. *Journal of Artificial Intelligence Research*, 30(1):133–179.
- Buchbinder, N., Jain, K., and Singh, M. (2010). Incentives in online auctions via linear programming. In *Proceedings of the 6th International Conference on Internet and Network Economics*, WINE’10, pages 106–117.
- Byers, J. W., Mitzenmacher, M., Potamias, M., and Zervas, G. (2011). A month in the life ofgroupon. *Computing Research Repository*, abs/1105.0903.
- Cai, K., Gerding, E., Mcburney, P., Niu, J., Parsons, S., and Phelps, S. (2009). Overview of cat: A market design competition version 2.0. Technical report, University of Liverpool.
- Cavallo, R. and Parkes, D. C. (2008). Efficient metadeliberation auctions. In *Proceedings of the 23rd National Conference on Artificial Intelligence - Volume 1*, AAAI’08, pages 50–56.
- Cavallo, R., Parkes, D. C., and Singh, S. (2009). Efficient mechanisms with dynamic populations and dynamic types. Technical report, University of Chicago Booth School of Business.
- Chrobak, M. (2008). Sigact news online algorithms column 13: 2007 - an offline perspective. *SIGACT News*, 39:96–121.
- Chu, L. Y. (2009). Truthful bundle/multiunit double auctions. *Management Science*, 55:1184–1198.
- Cliff, D. and Bruten, J. (1997). *Minimal-intelligence Agents for Bargaining Behaviors in Market-based Environments*. HP Laboratories Technical Report. Hewlett Packard Laboratories.
- Conitzer, V. and Sandholm, T. (2007). Incremental mechanism design. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, IJCAI’07, pages 1251–1256.

- Constantin, F. and Parkes, D. C. (2009). Self-correcting sampling-based dynamic multi-unit auctions. In *Proceedings of the 10th ACM Conference on Electronic Commerce, EC'09*, pages 89–98.
- Dash, R. K., Jennings, N. R., and Parkes, D. C. (2003). Computational-Mechanism design: A call to arms. *IEEE Intelligent Systems*, 18(6):40–47.
- Edelman, B., Jaffe, S., and Kominers, S. D. (2011). Togroupon or not togroupon: The profitability of deep discounts. Harvard Business School Working Papers 11-063, Harvard Business School.
- Erev, I. and Roth, A. E. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review*, 88(4):848–81.
- Friedman, D. and Rust, J. (1993). *The Double Auction Market: Institutions, Theories, And Evidence*. Westview Press.
- Galil, Z. (1986). Efficient algorithms for finding maximum matching in graphs. *ACM Computing Surveys*, 18:23–38.
- Gjerstad, S. and Dickhaut, J. (2001). Price formation in double auctions. In *E-Commerce Agents, Marketplace Solutions, Security Issues, and Supply and Demand*, pages 106–134.
- Gode, D. K. and Sunder, S. (1993). Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–37.
- Goldberg, A. V., Hartline, J. D., Karlin, A. R., Wright, A., and Saks, M. (2002). Competitive auctions. In *Games and Economic Behavior*, pages 72–81.
- Gonen, M., Gonen, R., and Pavlov, E. (2007). Generalized trade reduction mechanisms. In *Proceedings of the 8th ACM Conference on Electronic Commerce, EC'07*, pages 20–29.
- Groves, T. (1973). Incentives in teams. *Econometrica*, 41(4):617–31.
- Hajiaghayi, M. T. (2005). Online auctions with re-usable goods. In *Proceedings of the 6th ACM Conference on Electronic Commerce, EC'05*, pages 165–174.

- Hajiaghayi, M. T., Kleinberg, R., and Parkes, D. C. (2004). Adaptive limited-supply online auctions. In *Proceedings of the 5th ACM Conference on Electronic Commerce, EC'04*, pages 71–80.
- Honari, S., Ebadi, M., Foshati, A., Gomrokchi, M., Benatahr, J., and Khosravifar, B. (2009). Price estimation of persiancat market equilibrium. IJCAI Workshop on Trading Agent Design and Analysis (TADA).
- Hopcroft, J. E. and Karp, R. M. (1971). A $n^{5/2}$ algorithm for maximum matchings in bipartite. In *Proceedings of the 12th Annual Symposium on Switching and Automata Theory, SWAT'71*, pages 122–125.
- Huang, P., Scheller-wolf, A., and Sycara, K. (2002). Design of a multi-unit double auction e-market. *Computational Intelligence*, 18:596–617.
- Kalyanasundaram, B. and Pruhs, K. (1993). Online weighted matching. *Journal of Algorithms*, 14:478–488.
- Karp, R. M., Vazirani, U. V., and Vazirani, V. V. (1990). An optimal algorithm for on-line bipartite matching. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing, STOC'90*, pages 352–358.
- Kleinberg, R. (2005). A multiple-choice secretary algorithm with applications to online auctions. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA'05*, pages 630–631.
- Kozen, D. (1991). *The Design and Analysis of Algorithms*. Springer-Verlag.
- Mahdian, M. and Saberi, A. (2006). Multi-unit auctions with unknown supply. In *Proceedings of the 7th ACM Conference on Electronic Commerce, EC'06*, pages 243–249.
- McAfee, R. P. (1992). A dominant strategy double auction. *Journal of Economic Theory*, 56(2):434–450.
- Mehta, A., Saberi, A., Vazirani, U., and Vazirani, V. (2007). Adwords and generalized online matching. *Journal of the ACM*, 54.
- Myerson, R. B. (2008a). mechanism design. In Durlauf, S. N. and Blume, L. E., editors, *The New Palgrave Dictionary of Economics*. Palgrave Macmillan.
- Myerson, R. B. (2008b). revelation principle. In Durlauf, S. N. and Blume, L. E., editors, *The New Palgrave Dictionary of Economics*. Palgrave Macmillan.

- Myerson, R. B. and Satterthwaite, M. A. (1983). Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281.
- Nagle, T. and Holden, R. (2001). *The Strategy and Tactics of Pricing: a guide to profitable decision making*. Prentice Hall.
- Nisan, N. (2007). Introduction to mechanism design (for computer scientist). In *Algorithmic Game Theory*. Cambridge University Press.
- Nisan, N., Ronen, A., and Ronen, A. (1999). Algorithmic mechanism design. In *Games and Economic Behavior*, pages 129–140.
- Nisan, N., Roughgarden, T., Éva Tardos, and Vazirani, V. V. (2007). *Algorithmic Game Theory*. Cambridge University Press.
- Niu, J., Cai, K., and Parsons, S. (2010). A grey-box approach to automated mechanism design. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems - Volume 1, AAMAS'10*, pages 1473–1474.
- Niu, J., Cai, K., Parsons, S., Gerding, E., McBurney, P., Moyaux, T., Phelps, S., and Shield, D. (2008). Jcat: a platform for the tac market design competition. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS'08*, pages 1649–1650.
- Parkes, D. C. (2007). Online mechanisms. In *Algorithmic Game Theory*. Cambridge University Press.
- Parkes, D. C. and Duong, Q. (2007). An ironing-based approach to adaptive online mechanism design in single-valued domains. In *Proceedings of the 22nd National Conference on Artificial Intelligence - Volume 1, AAAI'07*, pages 94–101.
- Phelps, S., Mcburney, P., and Parsons, S. (2010). Evolutionary mechanism design: a review. *Autonomous Agents and Multi-Agent Systems*, 21:237–264.
- Robinson, E. R., McBurney, P., and Yao, X. (2009). How specialised are specialists generalisation properties of entries from the 2008 tac market design competition. IJCAI Workshop on Trading Agent Design and Analysis (TADA).
- Rothkopf, M. H. (2007). Thirteen reasons why the vickrey-clarke-groves process is not practical. *Operations Research*, 55(2):191–197.

- Sandholm, T. (2003). Making markets and democracy work: a story of incentives and computing. In *Proceedings of the 18th international joint conference on Artificial intelligence, IJCAI'03*, pages 1649–1671.
- Stavrogiannis, L. C. and Mitkas, P. A. (2009). Cat 2008 post-tournament evaluation: The mertacors perspective. IJCAI Workshop on Trading Agent Design and Analysis.
- Tarjan, R. E. (1983). *Data structures and network algorithms*. Society for Industrial and Applied Mathematics.
- Ünver, M. U. (2010). Dynamic kidney exchange. *Review of Economic Studies*, 77(1):372–414.
- Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37.
- Vytelingum, P., Vetsikas, I. A., Shi, B., and Jennings, N. R. (2008). Iamwildcat: The winning strategy for the tac market design competition. In *Proceeding of the 18th European Conference on Artificial Intelligence, ECAI'08*, pages 428–432.
- West, D. B. (2000). *Introduction to Graph Theory*. Prentice Hall, 2 edition.
- Wilson, R. (1985). Incentive Efficiency of Double Auctions. *Econometrica*, 53(5):1101–1115.
- Wurman, P. R., Walsh, W. E., and Wellman, M. P. (1998). Flexible double auctions for electronic commerce: theory and implementation. *Decision Support Systems*, 24:17–27.