

AUTOMATED NEGOTIATION FOR COMPLEX MULTI-AGENT RESOURCE ALLOCATION

A Dissertation Presented

by

BO AN

Submitted to the Graduate School of the
University of Massachusetts Amherst in partial fulfillment
of the requirements for the degree of

DOCTOR OF PHILOSOPHY

February 2011

Department of Computer Science

© Copyright by Bo An 2011

All Rights Reserved

AUTOMATED NEGOTIATION FOR COMPLEX MULTI-AGENT RESOURCE ALLOCATION

A Dissertation Presented

by

BO AN

Approved as to style and content by:

Victor Lesser, Chair

Jim Kurose, Member

Shlomo Zilberstein, Member

Michael Zink, Member

Andrew Barto, Department Chair
Department of Computer Science

ACKNOWLEDGMENTS

I take this opportunity to express my gratitude to all the people who have provided me with their help and without those help this work would not been possible. First and foremost, I want to thank my advisor Victor Lesser for his constant support and encouragement in all aspects of my research, career and even personal issues. Victor has been not only a great advisor but also a cordial friend and a person mentor. He is always enthusiastic in sharing his experience and knowledge with me. He has taught me how to do deep research, to find new research directions by discovering variations and generalizations, to deliver presentations, and to be confident. He gave me a lot of freedom in pursuing my broad interests and tried very hard to put me in close contact with the multiagent systems community. I am grateful for the opportunity to work with and to learn from such an extraordinary teacher.

I also especially want to thank the other members of my committee, Jim Kurose, Shlomo Zilberstein, and Michael Zink, for their helpful and insightful comments. Special thanks go to Jim Kurose for advising my synthesis project and Michael Zink for guiding my research on the cloud computing resource allocation problem in Chapter 6.

I gratefully thank the NSF CASA project (contract No. EEC-0313747) for providing me with funding over the last four years. Working for the NSF CASA project has taught me much about how to do research which has great practical impact and I am grateful for all the interactions that I have had with the great people there, including Sherief Abdallah, Michael Krainin, Eric Lyons, Dave Pepyne, Brenda Phillips, David Westbrook, and Michael Zink. I am also very grateful for the University of Massachusetts Amherst Graduate School Fellowship that partially supported me this past year.

I want to thank other co-authors who provided me valuable suggestions and guidance including Nicola Gatti for the work on bargaining theory, Kwang Mong Sim for the work on multi-resource negotiation, and David Irwin for the work on cloud computing resource allocation. I also want to thank Fred Douglass and Fan Ye for advising my intern research work at IBM research. I am indebted to many other researchers in multiagent systems and electronic commerce with whom I have had valuable discussions. I would like to particularly thank Edith Elkind, David Parkes, Tuomas Sandholm, Katia Sycara, Milind Tambe, Thanos Vasilakos, and Xiaoqin Zhang for their support toward my research work.

The Multi-Agent Systems Lab has been a family. I have benefited greatly from many technical discussions with Dan Corkill, Alan Carlin, Yoonheui Kim, He Luo, Hala Mostafa, Mark Sims, Huzaiifa Zafar, Chongjie Zhang, and many others. I also would like to thank Michele Roberts for helping with a variety of paperworks.

Before coming to UMASS, I had the good fortune to work with many excellent advisors. I would like to thank Daijie Cheng, Shuangqing Li, Chunyan Miao, Zhiqi Shen, Kwang Mong Sim, and Lianggui Tang for all their support and advice to build my career.

Finally, I want to thank my family for their unconditional love from the beginning. I want to thank my wife Jinghua for her love and support, especially for giving up everything in China before coming to the states. Thank you. I thank my daughter Amber. Your birth gives me much joy. This thesis is dedicated to them.

ABSTRACT

AUTOMATED NEGOTIATION FOR COMPLEX MULTI-AGENT RESOURCE ALLOCATION

FEBRUARY 2011

BO AN

B.Sc., CHONGQING UNIVERSITY

M.Sc., CHONGQING UNIVERSITY

M.Sc., UNIVERSITY OF MASSACHUSETTS AMHERST

Ph.D., UNIVERSITY OF MASSACHUSETTS AMHERST

Directed by: Professor Victor Lesser

The problem of constructing and analyzing systems of intelligent, autonomous agents is becoming more and more important. These agents may include people, physical robots, virtual humans, software programs acting on behalf of human beings, or sensors. In a large class of multi-agent scenarios, agents may have different capabilities, preferences, objectives, and constraints. Therefore, efficient allocation of resources among multiple agents is often difficult to achieve. Automated negotiation (bargaining) is the most widely used approach for multi-agent resource allocation and it has received increasing attention in the recent years. However, information uncertainty, existence of multiple contracting partners and competitors, agents' incentive to maximize individual utilities, and market dynamics make it difficult to calculate agents' rational equilibrium negotiation strategies and develop successful negotiation

agents behaving well in practice. To this end, this thesis is concerned with analyzing agents' rational behavior and developing negotiation strategies for a range of complex negotiation contexts.

First, we consider the problem of finding agents' rational strategies in bargaining with incomplete information. We focus on the principal alternating-offers finite horizon bargaining protocol with one-sided uncertainty regarding agents' reserve prices. We provide an algorithm based on the combination of game theoretic analysis and search techniques which finds agents' equilibrium in pure strategies when they exist. Our approach is sound, complete and, in principle, can be applied to other uncertainty settings. Simulation results show that there is at least one pure strategy sequential equilibrium in 99.7% of various scenarios. In addition, agents with equilibrium strategies achieved higher utilities than agents with heuristic strategies.

Next, we extend the alternating-offers protocol to handle concurrent negotiations in which each agent has multiple trading opportunities and faces market competition. We provide an algorithm based on backward induction to compute the subgame perfect equilibrium of concurrent negotiation. We observe that agents' bargaining power are affected by the proposing ordering and market competition and for a large subset of the space of the parameters, agents' equilibrium strategies depend on the values of a small number of parameters. We also extend our algorithm to find a pure strategy sequential equilibrium in concurrent negotiations where there is one-sided uncertainty regarding the reserve price of one agent.

Third, we present the design and implementation of agents that concurrently negotiate with other entities for acquiring multiple resources. Negotiation agents are designed to adjust 1) the number of tentative agreements and 2) the amount of concession they are willing to make in response to changing market conditions and negotiation situations. In our approach, agents utilize a time-dependent negotiation strategy in which the reserve price of each resource is dynamically determined by 1)

the likelihood that negotiation will not be successfully completed, 2) the expected agreement price of the resource, and 3) the expected number of final agreements. The negotiation deadline of each resource is determined by its relative scarcity. Since agents are permitted to decommit from agreements, a buyer may make more than one tentative agreement for each resource and the maximum number of tentative agreements is constrained by the market situation. Experimental results show that our negotiation strategy achieved significantly higher utilities than simpler strategies.

Finally, we consider the problem of allocating networked resources in dynamic environment, such as cloud computing platforms, where providers strategically price resources to maximize their utility. While numerous auction-based approaches have been proposed in the literature, our work explores an alternative approach where providers and consumers negotiate resource leasing contracts. We propose a distributed negotiation mechanism where agents negotiate over both a contract price and a decommitment penalty, which allows agents to decommit from contracts at a cost. We compare our approach experimentally, using representative scenarios and workloads, to both combinatorial auctions and the fixed-price model, and show that the negotiation model achieves a higher social welfare.

TABLE OF CONTENTS

	Page
ACKNOWLEDGMENTS	iv
ABSTRACT	vi
LIST OF TABLES	x
LIST OF FIGURES	xi
 CHAPTER	
1. INTRODUCTION	1
1.1 Introduction	1
1.2 Motivating Examples	3
1.2.1 Collaborating, Autonomous Stream Processing Systems (CLASP)	3
1.2.2 Global Environment for Network Innovations (GENI)	5
1.3 Automated Negotiation for Complex Resource Allocation Problems	10
1.3.1 Negotiation with Uncertainty	11
1.3.2 One-to-Many and Many-to-Many Negotiation	13
1.3.3 Multi-Resource Negotiation	15
1.3.4 Negotiation with Decommithment for Dynamic Resource Allocation in Cloud Computing	18
1.4 Main Contributions	21
1.5 Thesis Organization	22
 2. LITERATURE REVIEW ON AUTOMATED NEGOTIATION	 24
2.1 Negotiation as a Mechanism	25

2.1.1	Bargaining Mechanism Design	25
2.1.2	Alternating-offers protocol and its extensions	27
2.2	Equilibrium strategies in Strategic Bargaining Game	29
2.3	Designing Negotiation Agents	34
2.4	Auction Mechanism	40
2.5	Summary	42
3.	NEGOTIATION WITH UNCERTAIN RESERVE PRICES	43
3.1	Background	43
3.2	Bargaining with Complete Information	45
3.3	One-sided Uncertainty about Reserve Prices	49
3.3.1	Introducing Uncertainty	50
3.3.2	Existing Solutions in Literature	51
3.4	The Algorithm for Finding All Sequential Equilibria	55
3.4.1	High Level Idea of the Approach	55
3.4.2	Computation Reduction	59
3.4.3	The Algorithm	62
3.4.4	Off the Equilibrium Path Optimal Strategies	62
3.5	The Buyer's Equilibrium Offer	65
3.5.1	Pooling Choice Rule	66
3.5.2	Separating Choice Rule	68
3.6	The Seller's Equilibrium Offer	71
3.7	Equilibrium Existence	74
3.8	The Value of Equilibrium Strategies	83
3.8.1	Heuristic Based Strategies	84
3.8.2	Different Strategy Combination	85
3.8.3	Performance Measures and Results	86
3.8.4	The Value of Choosing Equilibrium strategies	88
3.8.5	Comparison of Social Welfare	90
3.9	Applications of the Approach	92
3.9.1	Bilateral Negotiation with Uncertain Discount Factor	92
3.9.2	Bilateral Multi-issue Negotiation with Uncertain Weights	95
3.10	Summary	99

4. EXTENDING ALTERNATING-OFFERS BARGAINING IN ONE-TO-MANY AND MANY-TO-MANY SETTINGS	102
4.1 One-to-Many Alternating-Offers Negotiation	102
4.1.1 Negotiation Mechanism	102
4.1.2 Agents' Equilibrium Strategies	104
4.1.3 Equilibrium Outcome Computation and Uncertain Information	111
4.2 Many-to-Many Alternating-Offers Negotiation	113
4.2.1 Negotiation Mechanism	113
4.2.2 Agents' Equilibrium Strategies	114
4.2.3 Considerations on Settings with Uncertain Information	120
4.3 Uncertainty about Reserve Prices	121
4.3.1 Introducing Uncertainty	121
4.3.2 Off the Equilibrium Path Optimal Strategies	123
4.3.3 Our Approach	125
4.3.4 The Buyer's Equilibrium Offer	125
4.3.4.1 Pooling Choice Rule	126
4.3.4.2 Separating Choice Rule	127
4.3.5 The Seller's Equilibrium Offer	129
4.4 Summary	132
5. STRATEGIC AGENTS FOR MULTI-RESOURCE NEGOTIATION	134
5.1 Background	134
5.2 Negotiation mechanism	138
5.2.1 Assumptions	138
5.2.2 The Negotiation Problem	140
5.2.3 The Negotiation Protocol	142
5.2.4 The Negotiation Strategy	143
5.3 Heuristics based Strategies	145
5.3.1 An overview of negotiation strategies	145
5.3.2 Different deadlines for different resources	148
5.3.3 Generating proposals	151
5.3.4 Maximum number of final agreements	158

5.4	Empirical evaluation and analysis	161
5.4.1	The methodology	161
5.4.1.1	Agent design	161
5.4.1.2	Experimental settings	163
5.4.1.3	Performance measure	165
5.4.1.4	Results	167
5.4.2	Observations	168
5.4.2.1	Observation 1	168
5.4.2.2	Observation 2	170
5.4.2.3	Observation 3	173
5.4.2.4	Observation 4	180
5.4.2.5	Observation 5	181
5.4.2.6	Sensitivity analysis	182
5.4.3	Analysis of properties	183
5.5	Summary	186
6.	NEGOTIATION WITH DECOMMITMENT FOR DYNAMIC RESOURCE ALLOCATION IN CLOUD COMPUTING	188
6.1	Negotiation Over Decommitment Penalty	189
6.1.1	Leveled-commitment contracting	191
6.1.1.1	Contracting game	192
6.1.1.2	Decommitting game	193
6.1.2	Optimal contracts	193
6.1.3	Efficiency of Negotiating Over Penalty in Two-player Game	195
6.2	Resource Allocation in GENI	198
6.3	The Negotiation Model	200
6.3.1	The Resource Allocation Problem	200
6.3.2	Negotiation Protocol	202
6.4	Buyers' Negotiation Strategy	205
6.5	Sellers' Negotiation Strategy	211
6.6	Empirical evaluation	213
6.6.1	Different Mechanisms	214

6.6.2	Experimental Settings and Measures	215
6.6.3	Results	217
6.6.3.1	Performance of the negotiation mechanism	217
6.6.3.2	Evaluating agents' negotiation strategies	222
6.6.3.3	Sensitivity analysis	223
6.7	Summary	225
7.	CONCLUSIONS AND FUTURE RESEARCH	226
7.1	Contributions	226
7.2	Future Research	228
	BIBLIOGRAPHY	232

LIST OF TABLES

Table	Page
3.1 Simulation parameters	80
3.2 Average number of sequential equilibria and percentage of games with sequential equilibria	81
3.3 Average computation time (in seconds)	83
5.1 Symbols used in this chapter	146
5.2 Experimental settings	163
5.3 Performance Measure	165
5.4 Experimental results for 10^6 runs (performance measures are defined in Table ??)	166
6.1 Symbols used in this chapter	202
6.2 Variables	214

LIST OF FIGURES

Figure	Page
1.1 Execution of a distributed job consisting 3 subjobs [27]. Owner Site 3 executes one subjob, and dispatches two subjobs to Site 1 and 2 for execution. Site 4 monitors Sites 1 and 2.....	5
1.2 Resource sharing with one consumer and one provider [1]	6
1.3 Resource sharing with multiple consumers and multiple providers	8
3.1 Backward induction construction with $RP_{\mathbf{b}} = 100$, $RP_{\mathbf{s}} = 0$, $\iota(0) = \mathbf{s}$, $\delta_{\mathbf{b}} = 0.75$, $\delta_{\mathbf{s}} = 0.8$, $T_{\mathbf{b}} = 10$, $T_{\mathbf{s}} = 11$; at each time point t the optimal offer $x^*(t)$ is marked; the dashed lines are isoutility curves.	48
3.2 Failure of the approach in [51, 52] with $T = 5$, $\iota(0) = \mathbf{s}$, $RP_{\mathbf{s}} = 10$, $RP_1 = 90$, $RP_2 = 70$, $\omega_{\mathbf{b}_1}^0 = 0.8$, $\omega_{\mathbf{b}_2}^0 = 0.2$, $\delta_{\mathbf{s}} = 0.7$, and $\delta_{\mathbf{b}} = 0.8$; agents' offers in complete information settings were also showed.	53
3.3 A high level illustration of our approach ($\iota(t) = \mathbf{s}$ and $ \Delta^0 > 1$).....	56
3.4 The buyer's different rates of concession	84
3.5 Buyer's average utility while using different strategies	86
3.6 Buyer's average utility and the negotiation deadline	87
3.7 Buyer's average utility and agents' discount factors.....	88
3.8 Seller's average utility while using different strategies	89
3.9 Seller's average utility and agents' discount factors	90
3.10 Social welfare for different strategy combinations.....	91
3.11 Social welfare and agents' discount factors	92

3.12	Social welfare and the negotiation deadline	93
3.13	Failure of the approach in [52, 51] with $T = 5$, $\iota(0) = \mathbf{s}$, $RP_s^1 = RP_s^2 = 0$, $RP_b^1 = RP_b^2 = 90$, $\delta_s = 0.5$, $\delta_b = 0.8$, $w_s = \langle 0.9, 0.1 \rangle$, $\{w_{b_1} = \langle 0.4, 0.6 \rangle, w_{b_2} = \langle 0.9, 0.1 \rangle\}$, $\omega_{b_1} = 0.9$, and $\omega_{b_2} = 0.1$; agents' offers in complete information settings were also showed.	97
4.1	Backward induction construction with $RP_b = 1$, $RP_{s_1} = 0$, $RP_{s_2} = 0.2$, $\delta_b = 0.8$, $\delta_{s_1} = 0.7$, $\delta_{s_2} = 0.8$, $T_b = 10$, $T_{s_1} = 11$, $T_{s_2} = 7$; at each time point t the optimal offer $x_a^*(t)$ that $\iota(t)$ can make is marked; the dashed lines are sellers' optimal offer if there is only one seller.	109
4.2	Backward induction construction. At each time t the optimal offer $x_B^*(t)$ or $x_S^*(t)$ is marked.	119
5.1	Buyer a 's multi-resource negotiation problem	140
5.2	Prediction accuracy of <i>HBAs</i>	171
5.3	Deadline and expected utility	174
5.4	Deadline and success rate	175
5.5	Number of resources and expected utility	176
5.6	Number of resources and success rate	177
5.7	Supply/demand ratio and expected utility	178
5.8	Supply/demand ratio and success rate	179
6.1	Efficiency comparison in two-player game (1)	197
6.2	Efficiency comparison in two-player game (2)	198
6.3	Finite state machine for the negotiation protocol	203
6.4	Social welfare and resource competition	216
6.5	Success rate and resource competition	217
6.6	Social welfare and number of resource to acquire	218

6.7	Success rate and number of resource to acquire	219
6.8	Social welfare and the flexibility of starting a task.....	220
6.9	Success rate and the flexibility of starting a task	221
6.10	Social welfare and negotiation time	222
6.11	Success rate and negotiation time	223

CHAPTER 1

INTRODUCTION

1.1 Introduction

The problem of resource allocation is ubiquitous in many diverse research fields such as economics, operations research, and computer science, and is relevant to a wide range of applications, e.g., electronic commerce, supply chain, sensor networks, web/grid service composition, workflow, and enterprise integration. In systems involving multiple autonomous agents, it is often necessary to decide how scarce resources should be allocated. The allocation of resources within a system of autonomous agents is a challenging and exciting area of research at the interface of computer science and economics. There are (at least) two different lines of work, depending on how decisions about allocations are made. In centralized approaches like combinatorial auctions [26, 42], agents simply report their preferences and wait for the final allocation to be made by the auctioneer or some other central entity. However, it is often impractical to adopt a centralized approach for resource allocation problems due to various constraints, e.g., computational overhead, communication constraints, privacy, and real-time requirements. In distributed approaches, allocations evolve in an asynchronous way, by means of local negotiations among agents.

Negotiation has been treated as a key approach for resource allocation problems and has been applied to e-commerce, manufacturing planning, sensor networks, cloud computing, and distributed vehicle routing [10, 75, 119]. Automated negotiation is an important research area bridging together economics, game theory, and artificial intelligence. It has received prominent attention in recent years [70] and its impor-

tance is widely acknowledged since intelligent agents that negotiate with each other on behalf of human users are expected to lead to more efficient negotiations [120]. A very common class of negotiation is bargaining.¹ It refers to a situation in which individual agents have the possibility of concluding a mutually beneficial agreement which could not be imposed without all individuals' approval. In intelligent agent systems like electronic markets, agents negotiate with each other through some form of negotiation protocol and negotiation capabilities for software agents are a central concern. Specifically, agents need to employ certain strategies to make negotiation decisions on behalf of the parties they represent with the aim of maximizing benefit for their users. One open question about automated negotiation is determining which strategy to employ, which is a complex decision making task because of the inherent uncertainty and dynamics of the situation.

In this thesis, we will be concerned with agents' negotiation strategies in dynamic and complex negotiation environments in which 1) agents have conflicting objectives and preferences; 2) agents need to acquire a set of resources; 3) agents have incomplete information about others; and 4) agents have multiple trading partners and trading competitors. The complex negotiation problem has not been studied deeply in the field. Our approaches will include both theoretic analysis and heuristic implementation. From a theoretical perspective, we analyze agents' rational strategies in bargaining games that more closely mirror the issues in complex negotiation than have previously been studied. Game theoretic analysis provides insights and theoretical foundations for developing negotiation agents. For more realistic complex dynamic bargaining games involving multiple agents, it is impractical to compute agents' rational strategies and we design heuristics based negotiation strategies by considering

¹Unless a specific distinction is drawn, we use the terms negotiation and bargaining interchangeably.

agents' constraints, contracting opportunities and market competition.² As part of this work, we will also focus on the use of decommitment penalties to handle the uncertainty present in dynamic and complex bargaining environments.

1.2 Motivating Examples

The focus of this thesis is developing new techniques for complex agent-mediated negotiation problems in which agents have different goals, preferences, constraints, and knowledge about others. To motivate the research from a practical perspective, this section describes two examples of application domains where the approaches developed in this thesis are needed.

1.2.1 Collaborating, Autonomous Stream Processing Systems (CLASP)

Collaborating, Autonomous Stream Processing Systems (CLASP) [27] is a middleware for cooperating data stream processing sites, which has been designed and prototyped in the context of System S project [69] within IBM Research to enable sophisticated stream processing. There are multiple sites running the System S software, each with their own administration and goals. Each site may only have limited processing capabilities, so cooperation among these sites can frequently be of mutual benefit and such cooperation enables such sites to increase the scale, breadth, depth, and reliability of analysis beyond that available within a single site [5].

In System S, a job is an execution unit that accomplishes certain work through stream analysis. A job takes the form of a processing graph, consisting of resources, i.e., data sources and processing elements (PEs), which are interconnected in a certain manner. These resources might be located at multiple different sites. Due to the potentially large numbers of data sources and PEs needed in complex jobs, and the

²Informally, a rational agent is an agent which has clear preferences, models uncertainty via expected values, and always chooses to perform the action that results in the optimal outcome for itself from among all feasible actions.

existence of functionally equivalent processing graphs, it is infeasible for human users to manually construct and identify the best alternative graph. System S has a planning component that can construct processing graphs automatically from high-level descriptions of desired results [108].

Figure 1.1 illustrates an example of a distributed job execution [27]. Site 3 is responsible for executing a job, which can be decomposed to three subjobs, each of which contains a *normal* job (for data processing) and multiple tunneling PE jobs (for data transportation). The Remote Execution Coordinator (REC) of owner site 3 executes the third subjob and dispatches two other subjobs to Sites 1 and 2 for execution. Site 4 monitors the execution status of Sites 1 and 2. The REC at the owner site maintains a subjob table about which subjobs are running at which other site. The table is used for recovery of subjobs on failed sites. The REC executing a subjob first parses its Job Description Language (JDL) to identify one normal job, and multiple tunneling PE jobs. One thread is launched to handle each of them. The thread customizes the JDL, such as assigning a host for each PE. Then it deploys the job through its local Job Management. For a source PE job, the REC needs to contact the local Tunneling Manager responsible for assigning the network address and port on which the source PE will be listening for incoming connections. It deploys the source PE job and reports the assigned network location to the REC at the owner site. For a sink PE job, the REC needs to query the REC of the owner site for the network location of the corresponding source PE. Then it configures and deploys the sink PE job.

Many resources needed in plans are accessed exclusively. In order for a site to reserve a limited resource from another site, it must establish an agreement with the other site, specifying the price of sharing the resource. Consider that a site receives a job. After planning [108], the site finds that using only its local resources, it cannot satisfy all resource requirements of the plan. Then, the site negotiates with




Figure 1.1. Execution of a distributed job consisting 3 subjobs [27]. Owner Site 3 executes one subjob, and dispatches two subjobs to Site 1 and 2 for execution. Site 4 monitors Sites 1 and 2.

other sites to acquire resources needed using its negotiation management component [5]. For each resource, there can be multiple providers and the site negotiates with different resource providers to construct agreements for these resources. The plan can be executed if and only if *all* resource requirements are satisfied. Therefore, while making a proposal to a trading partner for one resource, the site needs to consider the dynamically changing negotiation environments (e.g., the number of sites requiring the same resource) and the negotiation situation of other negotiations for the same resource and for other resources. It also has to consider the total price it will need to pay for all the resources needed to complete the job. Chapter 5 presents the design of negotiation strategies for a multi-resource allocation problem abstracted from this example.

1.2.2 Global Environment for Network Innovations (GENI)

Cloud computing platforms enable consumers to programmatically rent multiple types of Internet-accessible computing resources. In many cases, these platforms use recent advances in virtualization to make the resources appear to the consumer as raw

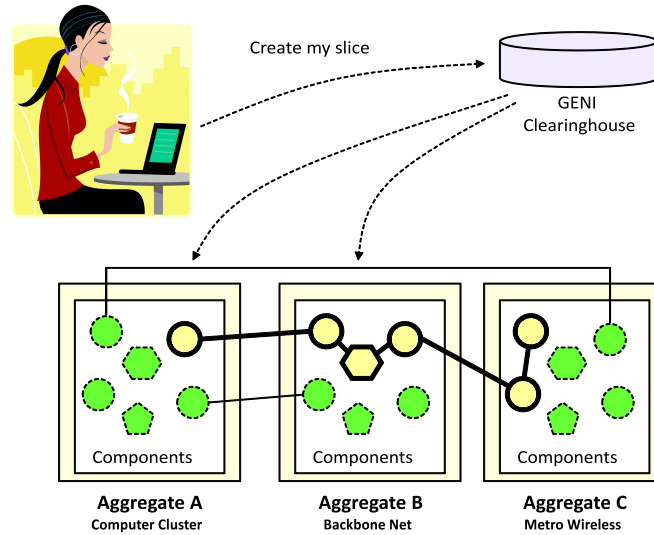


Figure 1.2. Resource sharing with one consumer and one provider [1]

hardware components, such as machines, storage block devices, sensors, or network links. For example, Amazon currently operates both the Elastic Compute Cloud (EC2) and the Elastic Block Store (EBS), where consumers programmatically rent virtual machines and block devices, respectively. As another example, the Global Environment for Network Innovations (GENI) project [2] is a recent NSF initiative that uses a similar paradigm but incorporates a wider range of hardware components, including not only machines and block devices, but also sensors, mobile devices, and the network links connecting them, from a wider range of providers, including universities and industry research labs.

GENI aims to provide a flexible and programmable shared experimental infrastructure for the investigation of future internet protocols and software. As explained in the GENI System Overview [1], one core concept for the suite of GENI infrastructure is Resource Sharing. That is, multiple researchers can simultaneously share the infrastructure and each experiment runs within its own, isolated slice created end-to-

end across the experiment's GENI resources. Furthermore, GENI experiments will be an interconnected set of reserved resources on platforms in diverse locations. Researchers will remotely discover, reserve, configure, program, debug, operate, manage, and teardown distributed systems established across parts of the GENI suite.

To illustrate the GENI's basic concepts regarding resource sharing, we show an example with a researcher who wishes to use GENI to perform an experiment [1]. A GENI Clearinghouse can be treated as a resource provider which can provide a wide range of resources. Most GENI resource components are not treated as isolated units. Instead they are parts of aggregates, which are collections of resources managed as a coherent whole. GENI contain many different kinds of aggregates. For example, the Clearinghouse in Figure 1.2 controls three aggregates: a computing cluster, a backbone network, and a metropolitan wireless network.

One core concept of GENI is virtualization in the sense that multiple researchers can simultaneously share the infrastructure. If the clearinghouse agrees to provide resources for the researcher, it will create a *slice* for the researcher. A slice is an empty container into which experiments can be instantiated and to which researchers and resources may be bound. The slice in Figure 1.2 extends across three aggregates: a computer cluster, a backbone network, and a metro wireless network. The resources within this slice are linked together to form a coherent virtual network in which an experiment can run. By virtualization, this slice will be isolated from other researchers' slices so that experiments running within the slice will behave consistently no matter what other researchers are doing within their own GENI slices. Once the slice is created, the researcher can download code into her slice, debugs, collects measurements, and iterates.

Generally there are multiple consumers acquiring resources and multiple clearinghouses which can provide resources (see Figure 1.3). Each clearinghouse may be operated by a private company, other US government agency, or indeed a separate




Figure 1.3. Resource sharing with multiple consumers and multiple providers

nation. For a consumer's resource requirement, there could be multiple clearinghouses which can satisfy the consumer's resource requirement separately. In this situation, the consumer has the opportunity to choose the clearinghouse with the lowest cost. It is also possible that a consumer's resource requirement cannot be satisfied by any single clearinghouse. In this situation, the consumer needs to acquire resources from multiple clearinghouses which may need to coordinate with each other to satisfy the consumer's resource requirements. After the consumer makes agreements with clearinghouses regarding resource sharing, a slice will be created over multiple clearinghouses.

Each consumer achieves some utility once its resource requirement is satisfied. A clearinghouse suffers a cost while providing resources. Given the existence of many resource consumers, the resources provided in the market is "limited". Therefore, the problem of allocating resources is an important issue. Since resource consumers and providers are always trying to maximize their own utilities, it is necessary to introduce some market mechanism to regulate the behaviors of resource consumers and providers: clearinghouses charge consumers and in turn consumers pay clear-

inghouses. Automated negotiation can be used for agents to make agreements on resource sharing. Chapter 6 discusses a distributed negotiation mechanism for the dynamic resource allocation problem in GENI.

The negotiation problem in the above two examples has the following features:

- Each agent has a negotiation deadline. A resource consumer's negotiation deadline is the time by which its job has to be executed. The presence of deadline indicates that an agent may need to make larger concessions when its deadline is approaching.
- All agents (including all resource consumers and resource providers) are selfish. That is, during negotiation, each agent chooses its negotiation strategy maximizing its (expected) utility. While agents are not cooperative, an agent often has no *optimal* strategy and we use the notion of an *equilibrium* strategy to define rational behavior of players, which jointly decide the outcome of the bargaining game. A strategy equilibrium is a profile of players' strategies so that no player could benefit by unilaterally deviating from its strategy in the profile, given that other players follow their strategies in the profile.
- Each agent has incomplete information about others. In a stream processing system, a site has incomplete information about other sites' cost of providing resources. Similarly, a resource provider does not know the exact reserve price of a site (the highest price the site is willing to pay) which desires its resource. Furthermore, resource supply and requirement in a stream processing system change dynamically.
- There could be multiple resource providers for a resource. To acquire a resource, a resource consumer can negotiate concurrently with all resource providers and make an agreement with the lowest price provider.

- A resource consumer also faces market competition from other resource consumers, which indicates that a negotiation agent needs to take the market situation into account to decide what is a necessary price to pay.
- An agent may need to acquire a set of resources and it gains nothing if it fails to get all the resources. Therefore, while making a proposal to a trading partner for one resource, the site needs to consider the dynamically changing negotiation environments (e.g., the number of sites requiring the same resource) and the negotiation situation of other negotiations for the same resource and for other resources, i.e., it is limited in what it can pay for the needed resources.
- In part of this work, we treat decommitment as a feature of negotiation problems. Since agents can choose to decommit from agreements, an agent may need to make more than one agreement for each resource. However, the buyer needs to pay much more by making more agreements. Thus, it is important to decide how many agreements to make. In some situations, it is also important to negotiate over both price and decommitment penalty.

The two examples demonstrate many of the issues that will be the focus of the thesis: negotiation where there is uncertainty about agents' types, negotiation involving multiple buyers and sellers, and negotiation where the buyer needs to acquire multiple resources.

1.3 Automated Negotiation for Complex Resource Allocation Problems

In designing automated negotiation agents for complex resource allocation problems we consider in this thesis, and in many others besides, there are a number of open problems and common issues that need to be dealt with. In addition, it is possible to identify a range of concepts and approaches that form a solid foundation for tackling

complex resource allocation problems. This thesis addresses important issues regarding automated negotiation for complex resource allocation problems. Specifically, this thesis investigates the following three questions:

- What are agents' equilibrium strategies in bilateral negotiation with uncertain reserve prices?
- How to compute agents' equilibrium strategies in concurrent one-to-many and many-to-many negotiation?
- When an agent needs to acquire multiple resources in a marketplace and it is allowed to decommit from an existing contract, how to make proposing decisions and decommitting decisions to achieve a high utility?
- How to develop an efficient negotiation model for resource allocation problems in dynamic markets such as cloud computing?

1.3.1 Negotiation with Uncertainty

The problem of finding agents' rational strategies in bargaining with incomplete information is well known to be challenging and there is no generally applicable algorithm [61]. We focus on finding agents' rational strategies in incomplete information bilateral bargaining. We consider the most common bargaining protocol, i.e., the Rubinstein's alternating-offers [111], which has been widely used in the bargaining theory literature, e.g., [61, 112, 117]. We analyze the situation with one-sided uncertain reserve prices and where agents have deadlines. This problem is customarily modeled as a Bayesian extensive-form game with infinite number of strategies as the price is a continuous value. The appropriate solution concept for such a class of game is *sequential equilibrium* [79], specifying a pair: a *system of beliefs* that prescribes how agents' beliefs must be updated during the game and *strategies* that prescribe how agents should act. In a sequential equilibrium there is a sort of circularity between

the belief system and strategies: strategies must be *sequentially rational* given the belief system and belief system must be *consistent* with respect to strategies.

The first area of focus is on the development of a novel algorithm to find a pure strategy sequential equilibrium in bilateral bargaining with multi-type uncertainty (Chapter 3). Our algorithm combines together game theoretic analysis with state space search techniques and it is sound and complete. Our approach is based on the following two observations: 1) with pure strategies, the buyer's possible choice rules regarding whether different buyer types behave in the same way or in different ways at a decision making point are finite, and 2) with pure strategies, the seller's possible beliefs regarding whether different buyer types will accept or reject its offer are finite. We employ a backward approach to find sequential equilibria in the context of a forward search process: to compute agents' equilibrium strategy at a continuation game with certain belief, we search forward to find agents' equilibria strategies in its continuation game with different beliefs and consider agents' all possible choice rules as well as belief update rules. At the same time, we derive theoretically the agents' optimal strategies by applying a Bayesian extension of backward induction and check equilibrium existence conditions.

In addition to developing the algorithm for computing all the sequential equilibria, we also empirically evaluate the performance of equilibrium strategies against some representative heuristic based strategies in the literature (e.g., [12, 48, 78, 124, 127]). Empirical results show that agents with equilibrium strategies achieved higher utilities than agents with heuristic based strategies. Furthermore, when both agents adopt the equilibrium strategies, they achieved the highest social welfare than that in all other strategy combinations.

1.3.2 One-to-Many and Many-to-Many Negotiation

In the bargaining theory literature, most work focuses on bilateral bargaining. A variety of negotiation aspects like information and outside options have been studied. One-to-many and many-to-many negotiations are also very important and widely exist in many application domains. For one-to-many negotiation, an auction is widely used and, for many-to-many negotiation, market mechanisms like matching or two-sided auction seem more intuitively appropriate. Even if an agent interacts with many agents, a common assumption in the literature is that an agent can pursue only one negotiation at a time. The result is that an agent may terminate a current negotiation in disagreement, in spite of possible gains from trade in order to pursue a more attractive outside alternative. Therefore, the presumption that an agent can pursue only one negotiation at a time appears to be restrictive. While there has been much experimental work (e.g., [97, 125]) on one-to-many and many-to-many negotiations in which an agent concurrently negotiates with multiple agents in discrete time, there is no game theoretic analysis of agents' strategic interactions in concurrent one-to-many and many-to-many negotiations.

The difference between negotiation and market mechanisms (e.g., auctions) is blurred with the arrival of the Internet and electronic commerce [72]. Negotiation has been treated as a key component of e-commerce and has been applied to e-commerce, manufacturing planning, and distributed vehicle routing. While an auction is the most widely implemented and discussed market mechanism, only recently the complex, multidimensional, and combinatorial auctions have gained the interest of researchers and foremost practitioners. Negotiations have been somewhat neglected as a possible market mechanism. The proliferation and acceptance of web and Internet technologies made the replacement of some negotiated transactions with auctions not only possible but also efficient. However, negotiation-based mechanisms still remain the preferred choice when the good and service attributes are ill-defined and there are

criteria other than price (e.g., reputation, trust, relation and future contracts) [64]. In addition, no third party like an auctioneer is needed in bargaining. Strategic agents may prefer bargaining as they can exploit other agents by using learning, collusion, and other bargaining techniques. In this work, we compared the efficiency of our model with some other mechanisms like auctions.

The second focus of this work is on analyzing agents' strategic behavior in one-to-many and many-to-many negotiations in which agents are negotiating with multiple trading partners and, at the same time, are facing competition from trading competitors (Chapter 4). The subgame perfect equilibrium for complete information setting is presented and equilibrium properties, such as uniqueness, are discussed. We analyze the reduction of computation in one-to-many settings and many-to-many settings. We also consider uncertainty about the reserve price of an agent while the reserve prices of other agents are common knowledge. We extend our approach for bilateral bargaining to search for sequential equilibrium when each agent is negotiating with multiple agents.

A central research topic in bargaining theory is understanding bargaining power, which is related to the relative abilities of agents in a situation to exert influence over each other. In bilateral bargaining, each agent's bargaining power is affected by its reserve price, patience attitude, deadline, etc. When many buyers and sellers are involved in negotiation, it is important to investigate how the market competition will affect agents' equilibrium bargaining strategies. With a large number of buyers and sellers, a single agent is unlikely to have much influence on the market equilibrium. Our analysis shows that both bargaining order and market competition affect agents' bargaining power. We show how an agent's bargaining power increases with the number of trading partners (agents of a different type) and decreases with the number of trading competitors.

1.3.3 Multi-Resource Negotiation

In electronic commerce markets where selfish agents behave individually, agents often have to acquire multiple resources in order to accomplish a high level task with each resource acquisition requiring negotiations with multiple resource providers. For example, in the CLASP [27], a site may need a set of resources to execute a job. Therefore, agents may need to engage in multiple negotiations. If the multiple negotiations are not all successful, consumers gain nothing. Such scenarios widely exist in practical applications. For example, a complex task may need several robots to work together and the absence of any of these robots results in the failure of the task. This is a simple form of multi-linked negotiation where the resources are independent but are interrelated. Resources are independent in the sense that there is no dependence between different resources, i.e., using one resource doesn't constrain how the other resources are used. However, from the perspective of the overall negotiation, resources are dependent as an agent's utility from the overall negotiation depends on obtaining overall agreements on all the resources. The negotiation problem we consider has the following three features:

1. When acquiring multiple resources, a consumer agent only knows the reserve price available for the entire set of resources, i.e., the highest price the agent can pay for all the resources, rather than the reserve price of each separate resource. In practice, given a plan and its resource requirements, an agent can easily decide the reserve price for all the resources in that plan based on the overall worth of the task. However, it is difficult (even impossible) for a resource consumer to understand how to set the reserve price for each separate resource. In fact, we show experimentally that it is undesirable to set a fixed reserve price for an individual resource prior to beginning negotiations.
2. Agents can decommit from tentative agreements at the cost of paying a penalty. Decommitment allows agents to profitably accommodate new tasks arriving

or new negotiation events. If these events make some existing contracts less profitable or infeasible for an agent, that agent can decommit from those contracts [121].

3. Negotiation agents are assumed to have incomplete information about other agents, for example, a buyer agent knows the distribution of the reserve price of a seller agent and the number of trading competitors. However, an agent's negotiation status (the set of proposals it has received) and negotiation strategy are its private information. For strategic reasons, a negotiation agent won't disclose such information during negotiation. During negotiation, negotiation agents can quit negotiation at any time, even without notifying their trading partners. When a buyer acquires multiple resources, it concurrently negotiates with sellers to reach agreements for all the resources.

Because resource providers and consumers may have different goals, preferences, interests, and policies, the problem of negotiating an optimal allocation of resources within a group of agents has been found to be intractable both in computation [45] and communication [47]. The multi-resource negotiation is even more complex due to decommitment. The multi-resource negotiation problem is different from multi-attribute negotiation in which negotiations are bilateral [49, 80, 82]. Sim and Shi [130] proposed a coordination strategy for multi-resource negotiation where an agent can negotiate with multiple agents as in this work. Each buyer in [130] knows the reserve price of each resource in advance and the buyer just needs to decide the concession strategy for each one-to-many negotiation for one resource. In contrast, each buyer in this work is assumed to only know the value of its high level task, i.e., the reserve price of all resources required for the high level task. Furthermore, a buyer in [130] only makes one tentative agreement but in this work, a buyer may make more than one tentative agreement. Nguyen and Jennings [97, 98] provide and evaluate a commitment model for concurrent negotiation. However, the maximum

number of tentative agreements is determined prior to negotiation. In our work, the maximum number of tentative agreements is determined by market situation and will change dynamically during negotiation. In addition, our work studies a multi-resource negotiation problem, rather than single resource negotiation as in [97, 98]. Furthermore, Nguyen and Jennings [97, 98] make very restrictive assumptions about agents' available information.

An agent's bargaining position in each round is determined by many factors such as market competition, negotiation deadlines, current agreement set, trading partners' proposals, and market dynamics. During each round of negotiation, an agent has to make decisions on how to proceed with each negotiation thread and there are many possible choices for each decision based on a variety of factors. Thus, it is difficult to construct an integrated framework in which all these factors are optimized concurrently. Rather than explicitly model those inter-dependent factors and then determining each agent's best decisions by an intractable combined optimization, the third area of focus (Chapter 5) tries to connect those inter-dependent factors indirectly and develops a set of heuristics to approximate agents' decision making during negotiation. The distinguishing feature of negotiation agents is that they are designed with the flexibility to adjust 1) the number of tentative agreements for each resource and 2) the amount of concession by reacting to i) changing market conditions, and ii) the current negotiation status of all concurrently negotiating threads. In our approach, agents utilize a time-dependent negotiation strategy in which the reserve price of each resource is dynamically determined by 1) the likelihood that negotiation will not be successful (*conflict probability*), 2) the *expected agreement price* of the resource, and 3) the *expected number of final agreements* given the set of tentative agreements made so far. The negotiation deadline of each resource is determined by its scarcity. A buyer agent can make more than one tentative agreement for each resource and the maximum number of tentative agreements is constrained by

the market situation in order to avoid the agent's making agreements more than necessary.

To evaluate the performance of negotiation agents, a simulation testbed consisting of a virtual e-Marketplace, a society of trading agents and a controller was implemented. Given that there is no existing negotiation agents dealing with our multi-resource negotiation problem, for comparison reason, we implemented three other types of buyers based on existing techniques for single resource negotiation and negotiation with decommitment. In the experiments, agents were subjected to different market densities, market types, deadlines, number of resources to acquire or sell, and supply/demand ratio of each resource. We use a number of performance measures including expected utility, success rate. Extensive stochastic simulations were carried out for all the combinations of market density, market type and other agents' characterizations. Experimental results show that the designed negotiation strategy achieved much better performance than other strategies.

1.3.4 Negotiation with Decommitment for Dynamic Resource Allocation in Cloud Computing

Cloud computing platforms enable consumers to programmatically rent multiple types of Internet-accessible computing resources. In many cases, these platforms use recent advances in virtualization to make the resources appear to the consumer as raw hardware components, such as machines, storage block devices, sensors, or network links. There are many reasons why market-oriented mechanisms are attractive for regulating resource supply and demand for these platforms. Amazon's goal is to make a profit by renting their resources to consumers for more than it costs to purchase and operate them. While GENI is initially operated as a non-profit platform, it allocates resources from multiple providers that dynamically donate and withdraw them, which makes centralized allocation difficult as the number of providers scales. Additionally,

market-oriented allocation mechanisms are attractive since they encourage providers to contribute resources to GENI in exchange for (real or virtual) currency that increases their access. Recent work has explored a variety of both system [55, 68] and market [19, 83] structures for resource allocation in market-oriented cloud computing platforms.

In this work, we focus on a general resource allocation problem that matches the characteristics of cloud computing platforms and their consumers. Namely, multiple self-interested agents supply or consume multiple types of resources, where 1) consumers dynamically enter and leave the market, 2) consumers have some bounded flexibility over when they require resources, and 3) a single provider cannot satisfy consumers' resource requirements. The first two characteristics are evident in current cloud platforms that are available to the general public, which use them to execute tasks that may or may not have hard deadlines. The motivation for 3) is natural for an infrastructure like GENI that allocates networked resources from multiple providers, and is also becoming more prevalent for profit-making enterprises like Amazon as competitors, such as RackSpace Cloud, become more prominent.

Given these characteristics, we consider the design of a market structure that allocates resources to their most *efficient* use. A straightforward approach would have all consumers submit both their resource requirements and bids to a single super agent that runs an auction, such as the well-known VCG auction [42], to allocate resources. Since VCG is not necessarily strategy-proof in dynamic settings, this approach does not necessarily result in the most efficient usage [101]. While efficient online mechanisms have been proposed for dynamic environments, they only work in constrained settings and often rely on strong assumptions about agents' knowledge [101]. Further, finding an auctioneer that selfish agents will trust and comply with is difficult. Alternatively, each consumer could run the VCG auction separately, but a provider may not truthfully report its information due to the existence of other auctions.

In this work, we present a negotiation mechanism in which agents make contracts for resource leases, which bind a set of resources from a provider to a consumer for a fixed time interval. To accommodate the highly dynamic nature of cloud computing platforms, we introduce a negotiation mechanism where an agent is able to decommit from a contract by paying a penalty to the other contract party. Thus, an agent may find it advantageous to decommit from existing contracts. Rather than setting decommitment penalties exogenously, we consider the role of negotiation in deciding decommitment penalties, where agents concurrently negotiate over both the contract price and the amount of decommitment penalty. We propose negotiating simultaneously over contract prices and decommitment penalties since it is difficult for system designers to decide the *optimal* contract prices and decommitment penalties that maximize the social welfare in dynamic environments involving multiple agents. We show that allowing decommitment improves the efficiency of the resource allocation mechanism.

Negotiation with uncertainty is both the most challenging problem in the negotiation literature [61], and is key to successful application of negotiation to real problems such as cloud computing. The literature provides a limited number of closed form results with narrow uncertainty settings using bilateral bargaining that considers only one type of uncertainty, such as a negotiation deadline [61] or reserve price [6]. In contrast, we consider negotiation between multiple agents in dynamic environments where there are multiple types of uncertainty that increases the difficulty of computing agents' rational equilibrium strategies. As a result, we bound agents' rationality and design negotiation strategies for them following the *negotiation decision functions* paradigm [11, 48]. Our negotiation problem is complex due to market dynamics, uncertainty, multiple contracting opportunities, resource competition, and decommitment. As a result, constructing an integrated framework for each agent that optimizes these factors concurrently is difficult. Rather than explicitly model

these inter-dependent factors and determine each agent’s best decisions through an intractable combined optimization, we connect these inter-dependent factors indirectly and develop a set of heuristics to approximate agents’ decision-making during negotiation. The distinguishing characteristic of our negotiation agents is their flexibility to adjust their decisions, such as making offers, by reacting to changing negotiation status, while also considering the time constraints, resource competition, and resource cost.

We evaluate our negotiation mechanism on a simulation testbed against two well-known mechanisms—combinatorial auctions and Amazon’s fixed-price model. Experimental results show that our negotiation mechanism achieves a higher social welfare than either mechanism in a wide range of scenarios. Further, we show that setting penalties through negotiation achieves a higher social welfare than exogenous mechanisms for setting penalties.

1.4 Main Contributions

The work described in this thesis makes a number of important contributions to the state of the art in the area of agent mediated negotiation by looking at more complex bargaining problems from both theoretical and heuristic perspectives. The contributions of this work can be summarized as follows:

- We present a novel algorithm to find a pure strategy sequential equilibrium in bilateral bargaining with multi-type uncertainty [6]. Our algorithm goes beyond existing algorithms dealing with complete information settings. Our approach is not specific to an application and it can be applied to other uncertainty settings, e.g., bilateral bargaining with uncertain discount factors, multi-issue negotiation with uncertain weight functions [51], and sequential auction (potentially over multiple goods).

- We extend the alternating-offers protocol to handle multiple trading opportunities and market competition [7]. We provide an algorithm based on backward induction to compute the subgame perfect equilibrium of concurrent one-to-many negotiation and many-to-many negotiation. There is no existing work on analyzing agents' equilibrium strategies in concurrent negotiation in markets.
- We present the design, implementation, and experimental evaluation of negotiation agents that negotiate for multiple resources where agents don't know the reserve price of each resource and are allowed to decommit from existing agreements [11, 12]. Existing work only considers single resource negotiation and often make unrealistic assumptions about agents' knowledge.
- We propose a distributed negotiation mechanism for the problem of allocating networked resources in dynamic environment, such as cloud computing platforms. In our approach, providers and consumers automatically negotiate resource leasing contracts as well as decommitment penalties. Experimental results show the advantage of the negotiation model over different combinatorial auction mechanisms and Amazon's fixed price model. This is the first work that shows the importance of negotiation over decommitment penalties.

1.5 Thesis Organization

The rest of this thesis is structured in the following manner: In Chapter 2, we discuss related research on automated negotiation, including both theoretic work and empirical work. Subsequently, in Chapter 3, we present an algorithm for finding agents' equilibrium in pure strategies in bargaining with one-side uncertainty about agents' reserve prices. Next, in Chapter 4, we analyze agents' equilibrium strategies in one-to-many and many-to-many negotiations. After that, in Chapter 5 and Chapter 6, we consider more practical multi-agent resource allocation problems. In

Chapter 5, we present the design and implementation of agents that concurrently negotiate with other entities for acquiring multiple resources as the negotiation problem in Section 1.2.1. In Chapter 6, we discuss the distributed negotiation mechanism for the problem of allocating networked resources in cloud computing platforms such as the GENI platform in Section 1.2.2. We finally summarize the contributions in this thesis and outline future directions In Chapter 7.

CHAPTER 2

LITERATURE REVIEW ON AUTOMATED NEGOTIATION

Automated negotiation is an important research area bridging together economics, game theory, and artificial intelligence. Bargaining (or negotiation) refers to a situation in which individual agents have the possibility of concluding a mutually beneficial agreement which could not be imposed without all individuals' approval. A bargaining theory is an exploration of the relation between the outcome of bargaining and the characteristics of the situation. Cooperative bargaining theory (axiomatic approach) initiated by Nash [95] is concerned with the outcome of bargaining given the list of properties (e.g., stability, fairness) the outcomes are required to satisfy. In the non-cooperative bargaining theory (strategic approach), the outcome is an equilibrium of an explicit model of the bargaining process. The strategic bargaining has received more attention following Rubinstein's path-breaking work [111]. In this chapter we provide an extensive literature review on the research of non-cooperative negotiation in the fields of economics and artificial intelligence.¹ The research in the economics community mainly focuses on computing agents' equilibrium strategies and the research in the AI part contributes to the development of software agents which negotiate on behalf of their users in realistic environments in which it is often impossible to compute agents' equilibrium strategies. We also discuss some market mechanisms for resource allocation problems.

¹An interested reader can refer to [70, 77, 107] for further discussions.

2.1 Negotiation as a Mechanism

Automated negotiation research can be considered to deal with three broad topics: negotiation protocols, negotiation objects and strategies. According to [70], a negotiation protocol is a set of rules that govern the interaction which cover the permissible types of participants (e.g., the negotiators and any relevant third parties), the negotiation states (e.g., accepting bids, negotiation closed), the events that cause negotiation states to change (e.g., no more bidders, bid accepted) and the valid actions of the participants in particular states (e.g., which messages can be sent by whom, to whom, at what stage). This section discusses a variety of negotiation protocols and we start with the formal bargaining mechanism design in the mechanism design literature.

2.1.1 Bargaining Mechanism Design

Bargaining mechanism design generally focuses on *bilateral* monopoly, in which a buyer and a seller are bargaining over the price of an object (e.g., a good). Myerson-Satterthwaite theorem [94] is one of the most remarkable negative results in economics. Informally, Myerson-Satterthwaite theorem says that there is no efficient way for two parties to trade a good when they each have secret and probabilistically varying valuations for it, without the risk of forcing one party to trade at a loss. Myerson and Satterthwaite analyze bargaining as a *static direct revelation* game in which each player reports its type to a third party, and the third party chooses whether the object is transferred, and how much the buyer must pay. Chatterjee and Samuelson [36] analyze a strategic game in which both players make offers simultaneously, and the trade occurs at a price between the two offers if the seller's offer is less than the buyer's offer. This game is closely related to the direct revelation game since it is static. Moreover, it can be shown that for a particular class of examples, the

simultaneous-offers game implements the direct revelation game in which the outcome functions are chosen to maximize the players' *ex ante* utility.

It is unrealistic to use a bargaining mechanism that forces agents to walk away from known positive gains from a potential trade since such mechanisms violate a broad interpretation of sequential rationality [40]. Cramton [40] examined the bargaining problem as a sequential direct revelation game, focusing on both the role of incomplete information and sequential aspects of bargaining. The difference between the static direct revelation game in [94] and the sequential direct revelation game in [40] is that in the sequential game, the outcome functions not only determine the probability and terms of trade, but also dictate when trade is to take place. In the static game trade may occur only at time zero whereas in the sequential game trade may occur at different times depending on the players' reports of their private information. Analyzing sequential bargaining mechanisms enable one to infer what the players' learning process is over time and to study what bargaining outcomes are possible when the bargainers are unable to make binding agreements.

Efficient bargaining mechanisms strongly depend on the bargaining settings characterized by agents' preferences and knowledge. Athey and Segal [16] consider the problem of allocating a good between two players in each period of an infinite-horizon game. The players' valuations in each period are private information, and the valuations change over time following a first-order Markov process. They analyze conditions under which there exists an efficient, Bayesian incentive-compatible, individually rational, budget-balanced mechanism, when the mechanism designer has commitment power.²

²A mechanism is Bayesian incentive-compatible if telling truth is a Bayesian-Nash equilibrium of the game induced by the mechanism. A mechanism is individually rational if an agent can always achieve as much expected utility from participation as without participation. A mechanism is budget-balanced if there are no net transfers out of the system or into the system.

Mechanism design is a powerful theory for studying incentive problems in bargaining. We are able to characterize the set of attainable outcomes and determine optimal or efficient trading mechanisms. However, mechanism design has a number of weaknesses [18]. First, the mechanisms may depend on the traders' beliefs and utility functions, which are assumed to be common knowledge. In addition, it is difficult to find a "satisfactory" bargaining mechanisms [94]. In practice, bargainers use simple negotiation protocols (e.g., the most widely used alternating-offers protocol [111]) that do not depend on agents' beliefs or utility functions. Given a negotiation protocol, the focus is then on analyzing agents' equilibrium strategies in the strategic bargaining game. An important distinction between direct revelation games and strategic games is that the direct revelation game does not explicitly model the process of bargaining. The sequence of offers and replies that eventually leads to an outcome is not studied in the direct revelation game as it is in strategic games.

2.1.2 Alternating-offers protocol and its extensions

The most widely used bargaining protocol in strategic bargaining games is the alternating-offers protocol, which was pioneered by Stahl [132] and Rubinstein [111] in a setting with complete information. The alternating-offers game represents a very general bargaining rule: at any time, a bargainer may make a new offer or accept the most recent offer of its opponent. The alternating-offers protocol captures the most important features of bargaining: bargaining consists of a sequence of offers and decisions to accept or reject these offers. The alternating-offers protocol has been widely used in the bargaining theory literature, e.g., [61, 112, 117], just to name a few.

The original alternating-offers protocol is designed for the simple discrete time bilateral single-issue negotiation and the allowed actions include *offer* and *accept*. The alternating-offers protocol has been extended in a variety of ways to handle

more complex negotiation situations, e.g., deadline, one-to-many negotiation, and decommitment. In realistic applications, agents often face deadlines and the action *quit* allows an agent to quit a negotiation before its deadline approaches. In addition to bilateral negotiation, one-to-many and many-to-many negotiations are also very important and widely exist in many application domains like e-commerce as well as in human society [5, 97, 98]. In automated negotiation systems for self-interested agents, contracts have traditionally been binding and do not allow agents to efficiently deal with future events in the environment. Sandholm and Lesser [121] proposed leveled-commitment contracts which allow an agent to be freed from an existing contract at the cost of simply paying a penalty to the other contract party. A self-interested agent will be reluctant to decommit because the other contract party might decommit, in which case the former agent gets freed from the contract, does not incur a penalty, and collects a penalty from the other party. Despite such strategic decommitting, leveled-commitment increases the expected payoffs of all contract parties and can enable deals that are impossible under full commitment [121]. Negotiation with decommitment has been applied in a variety of applications [5, 97, 98].

The contract net protocol (CNP) [131] is a simple negotiation protocol for distributed problem-solving based on the notion of call for bids on markets. The original CNP protocol is for cooperative problem solving and it has a number of limitations. For example, in the original CNP model, a contractor can only respond to bids sequentially. However, in a multi-agent system, several managers may concurrently call for bids and it is important to give each contractor the opportunity to concurrently negotiate with multiple managers and optimize its utility. In addition, there is no counter-proposing in the CNP model. The original CNP protocol has been extended in different applications. In the work on TRACONET [115, 118], a formal model based on marginal cost calculation was proposed for bounded rational self-interested agents to make announcing, bidding and awarding decisions. In early CNP imple-

mentations, tasks were negotiated one at a time, which is insufficient. Sandholm [113] analyzed task reallocation where individually rational agents contract tasks among themselves based on marginal costs and propose different contract types to facilitate negotiation. Aknine *et al.* [4] proposed an extended version of CNP to support concurrent negotiation processed for the task contractor service provider. Later, Dang and Huhns [43] extended the model to allow counter-proposing.

We consider both bilateral negotiation and concurrent negotiation. We use the most widely used alternating-offers protocol to study bilateral bargaining. To accommodate one-to-many negotiation, we extend the alternating-offers protocol by introducing another action *confirm* to avoid agents' non-reasonable behaviors as in the ADEPT (Advanced Decision Environment for Process Tasks) multi-agent architecture [71]. If one seller \mathbf{s} accepts an offer from a buyer \mathbf{b} , buyer \mathbf{b} needs to confirm the acceptance to reach an agreement. Notice that, in absence of the action *confirm*, if buyer \mathbf{b} makes offers to multiple sellers and all these accept, buyer \mathbf{b} must buy multiple items. In presence of the action *confirm*, buyer \mathbf{b} is in the position to choose only one contract.

2.2 Equilibrium strategies in Strategic Bargaining Game

Strategic bargaining theory uses the notion of an equilibrium strategy to define rational behavior of bargaining agents, which jointly decide the outcome of a bargaining game. A strategy equilibrium is a profile of players' strategies so that no player could benefit by unilaterally deviating from its strategy in the profile, given that other players follow their strategies in the profile.

A game is with complete information if the preference information of a player is known to all other players, otherwise it is a game with incomplete information. When both agents have complete information about each other, the appropriate solution concepts are Nash equilibria for one-shot bargaining games and subgame

perfect equilibria (SPE) for multi-stage bargaining games. The strategies chosen by all players are said to be in Nash equilibrium if no player has anything to gain by changing only his or her own strategy unilaterally. A subgame perfect equilibrium refines the Nash equilibria in dynamic games. A strategy profile is a subgame perfect equilibrium if it represents a Nash equilibrium of every subgame of the original game. Rubinstein [111] studies the alternating-offers game with infinite horizon. It describes two-person bargaining as an extensive game with perfect information in which the players alternate offers. A key assumption is that the players are impatient. The main result gives conditions under which the game has a unique subgame perfect equilibrium and characterizes this equilibrium.

For a dynamic bargaining game with incomplete information, the appropriate solution concept for such a class of game is *sequential equilibrium* [58], specifying a pair: a *system of beliefs* that prescribes how agents' beliefs must be updated during the game and *strategies* that prescribe how agents should act. A belief gives, for each information set of the game belonging to the player, a probability distribution on the nodes in the information set. In a sequential equilibrium there is a sort of circularity between the belief system and strategies: strategies must be *sequentially rational* given the belief system and belief system must be *consistent* with respect to strategies. The study of bargaining with uncertain information is well known to be a challenging problem due to this circularity and there is no generally applicable algorithm for such problem in the literature. Operations research inspired algorithms such as Miltersen-Sorensen [92] work only on games with finite number of strategies, and therefore cannot be applied to bargaining in which each agent's strategy space is generally continuous. Enumeration based methods were used in [104] to compute Nash equilibria. They enumerate the agents' strategy supports. This approach cannot be applied in bargaining problems since the number of strategies of each agent is infinite. Several attempts to extend the backward induction method [58] have been

tried, but they work for very restrictive cases. This is because in the computation of the equilibrium they break down the circularity between strategies and the belief system. For example, Fatima *et al.* [51, 52] present an algorithm to produce equilibrium strategies in multi-issue bargaining with uncertain reserve prices. By exploiting backward induction, their algorithm searches agents' strategy space from the deadline to the beginning of negotiation with the initial beliefs. Once the optimal strategies at the beginning of negotiation have been found, the system of beliefs are designed to be consistent with them. However, the optimization in their approach is myopic since it did not take into account its information effects. As a result, the strategies found by their approach are not guaranteed to be sequentially rational given the designed system of beliefs [61]. We will discuss this further in Chapter 3.

The microeconomic literature provides a number of closed form results with very narrow uncertainty settings, e.g., uncertainty regarding deadlines, reserve prices, and discounting factors. Rubinstein [112] considered bilateral infinite horizon bargaining with uncertainty over two possible discount factors. Sandholm and Vulkan [117] consider a continuous time bilateral bargaining with uncertainty deadlines and show that the only sequential equilibrium outcome is one where the agents wait until the first deadline, at which point that agent concedes everything to the other. In other words, bargaining game is a “waiting” game. Gatti *et al.* [61] provided an algorithm to compute agents' equilibrium strategies in bilateral bargaining with one-sided uncertain deadlines. In bargaining models with incomplete information, there could be multiple sequential equilibria (e.g., [112]) and most of these equilibria are supported by optimistic conjectures by the uninformed player. Bikhchandani [23] discussed how to eliminate some equilibria by placing restrictions on beliefs off equilibrium paths. The *only* known result about bargaining with uncertain reserve prices is due to Chatterjee and Samuelson [34, 35] where they studied bilateral infinite horizon bargaining with two-type uncertainty over the reservation values. The absence of agents' deadlines

makes these two results nonapplicable to the situation we study. In this work, we present a novel algorithm to compute sequential equilibrium strategies for bilateral finite horizon bargaining with uncertain reserve prices. Our approach can be applied to many other bargaining settings, like the multi-issue negotiation considered in [51].

Another challenging problem in bargaining theory is multi-issue negotiation, which is more complex and challenging than a single-issue negotiation (an interested reader can refer to [81] for a more detailed review). With multiple issues, agents need to decide the *negotiation procedure* and *agreement implementation*. There usually exist different types of negotiation procedures [51] like *package (simultaneous) deal*, *separate negotiation*, and *sequential negotiation*. Package deal means two agents negotiate a complete package on all issues simultaneously. Separate negotiation means agents negotiate each issue separately (independently & simultaneously). For sequential negotiation in which two agents negotiate issue by issue sequentially, agents need to decide a negotiation *agenda* (order of negotiation issues) [50]. There are generally two ways to decide the negotiation agenda: endogenous, i.e., allow the bargainers to decide which issue they will negotiate next during the process of negotiation, and exogenous agendas, i.e., fix the agenda exogenously as part of the negotiation procedure. For agreement implementation, there can be two types: sequential and simultaneous. Sequential implementation means the agreement on each issue is implemented once it is reached, while simultaneous implementation is that agreements are implemented together when all issues are settled.

There are two important questions regarding multi-issue negotiation. The first question is determining the best procedure in terms of efficiency and the optimal procedures in different settings are different. Busch and Horstmann [30] show that if agents are heterogeneous, agents might have conflicting favors on the procedures. Lang and Rosenthal [84] argue that the package deal is better if agents' payoff functions are concave. When there is a risk of breakdown due to agents' deadline, it is

found that in a noncooperative bargaining model with alternating offers and time preferences, the timing of issues (the agenda) matters and simultaneous bargaining over “packages” should be a prevailing phenomenon [67]. Fatima *et al.* [51] study different procedures for bilateral multi-issue negotiation and show that the package deal is the optimal procedure. For incomplete information multi-issue negotiation, it is necessary to consider the signaling factor. Bac and Raff [20] study a case with two simultaneous and identical pies where agents can either choose sequential negotiation with sequential implementation or simultaneous negotiation with simultaneous implementation. They show that as long as there is incomplete information about bargaining strength, players may engage in issue-by-issue negotiations even if 1) the issues are perfect substitutes and players are only concerned with maximizing their gains from settling the complete set of issues, and 2) there are no transaction costs involved in negotiating a complete package. Busch and Horstmann [31] show that issue-by-issue bargaining arises endogenously as part of a separating equilibrium in which agenda choice is used to signal bargaining strength.

The other important question regarding multi-issue negotiation is selecting the negotiation agenda while using the issue-by-issue approach. Flamini [54] shows that there is a Pareto superior agenda among the issue-by-issue procedures. Fershtman [53] defines two different agendas with simultaneous implementation and shows that 1) when agents have identical preferences, the highest payoff under the big pie first agenda is higher than that under the small pie first agenda and 2) when agents have conflicting preferences, they prefer the first issue negotiated to be least important to themselves but most important to the opponent. In and Serrano [66] show that restricting agendas yields multiplicity of equilibrium outcomes because it creates strong forms of non-concave payoff frontiers. Busch and Horstmann [29] explore how bargaining conflicts and procedures interact to determine players’ bargaining costs in multi-issue bargaining settings. They show that when bargaining frictions take the form

of discounting and agreements are implemented as they are reached, issue-by-issue negotiation can generate bargaining costs different from those that occur if all issues are bargained simultaneously.

While the bargaining theory literature mainly focuses on bilateral bargaining, it also considers an agent's other contracting opportunities in terms of "outside options". That is, a negotiating agent can exit the current negotiation and negotiate with another trading partner. The outside option strategic factor has been explored in different ways. While Shaked and Sutton [123] look at fixed exogenous outside options in complete information settings, Fudenberg *et al.* [59] consider incomplete information bargaining with outside opportunities. Muthoo [93] studies a model of the situation in which two players are bargaining face-to-face over the partition of a unit size cake and, moreover, one of the players can choose to temporarily leave the negotiating table to search for an outside option. A main conclusion is that the equilibrium outcome does not depend on whether a bargainer is allowed (within the game form) to choose to return to the negotiating table to resume bargaining after having searched for some finite time. There are also some work (e.g., [33, 60]) on modeling outside option as a sequential search process, where an agent can choose to search for other offers and return to bargaining at any time. The search policy and bargaining strategies are related due to the search cost.

2.3 Designing Negotiation Agents

Research on negotiation in the economics field considers relatively simple bargaining scenarios, e.g., there are only two agents, each agent has knowledge about other agents' preference and goals. However, in realistic applications, agents often have high level tasks with complex structures in dynamic uncertain environments and it is often impossible to compute agents' equilibrium strategies. In addition, game theoretic analysis often makes strong assumptions about agents' knowledge, which limits

the practical applicability of game theoretic results. Furthermore, game theoretic solutions in which agents are assumed to be fully rational cannot be applied to realistic negotiation problems as, in practice, it's not reasonable to assume agents' full rationality. In contrast, agents adopting AI approaches often have bounded rationality and make "satisfying" decisions based on heuristics. Research in economics and AI have different methodologies and concerns and their contributions complement each other. Research in economics provides insights and theoretical foundations for designing good heuristics, and heuristic approaches provide approximate solutions for realistic negotiations problems.

Negotiation strategy: Heuristic search has been widely used by bounded rational agents to find approximate solutions. To build more flexible and sophisticated negotiation agents, Faratin *et al.* [48, 71] devised a negotiation model that defines a range of Negotiation Decision Functions (*NDFs*) for generating (counter-)proposals based on time, resource, and behaviors of negotiators. In the time-dependent tactics, an agent submits offers that change monotonically from the minimum (best) to the maximum (worst) of the deal that she can agree on, and the rate of change depends on time. There are different time-dependent tactics depending on the changing rate of offers. The pressure of deadline has been widely studied. For example, Kraus *et al.* [78] proposed a strategic model in which the passage of time was taken into account. It has been shown that if agents use sequential equilibrium strategies, negotiation will end rapidly. The resource-dependent tactics are similar to the time-dependent ones in which time is the sole considered resource. The resource-dependent tactics are modeled in the same way as the time-dependent ones by using the same functions. The difference is that the resource-dependent tactics either, 1) have dynamic value of the maximum available resource, or 2) make the changing rate function depend on an estimation of the amount of a particular resource. The behavior-dependent tactics compute the next offer based on the previous attitude of the negotiation opponent.

These tactics are especially important in cooperative problem solving negotiation settings, or integrative negotiations, by allowing agents to consider the other agents' behavior. Sim *et al.* [124, 127] consider other factors, such as competition, trading alternatives, and differences of negotiators, and propose *market-driven agents* (MDAs) which can make minimally sufficient concessions. A game theoretical analysis of this approach [125] shows that the strategies of MDAs are in sequential equilibrium and market equilibrium for some specific bilateral and multilateral negotiations. Like MDAs, our negotiation agents for multi-resource negotiation make negotiation decisions taking into account market dynamics and negotiation status of all negotiation threads for all resources.

Multi-issue negotiation: There are two different definitions of a negotiation *issue* in the literature. In papers like [49, 80, 82], an issue is an attribute (e.g., price, quality, delivery time) of a resource. In this case, multi-issue negotiation is bilateral. An issue can also be treated as a resource as in [51, 128, 130] and in this case, a buyer can negotiate with multiple sellers for each resource. If a seller has multiple resources, such multi-resource negotiation could be bilateral and each resource can be treated as an attribute. Multi-issue negotiation is more complex and challenging than single-issue negotiation as the solution space is multi-dimensional and it's often difficult to reach a *Pareto-efficient* solution [81]. Almost all the work on multi-issue negotiation focuses on bilateral negotiation and a variety of learning and searching methods are used, e.g., case-based reasoning [133], similarity criteria based search [49], decentralized search [80, 82]. Klein *et al.* [73] propose a simulated annealing based approach appropriate for negotiating such complex contracts that achieves near-optimal social welfare for negotiations with binary issue dependencies. Different from related work on bilateral multi-issue negotiation, this work studies multi-resource negotiation where resources are provided by multiple agents and thus an agent is negotiating with multiple trading partners.

One-to-many and many-to-many negotiation: In many situations, an agent has an opportunity to make an agreement with more than one trading partners. An agent may also face the competition from agents of the same type, e.g., a buyer in negotiation faces competition from other buyers. Even if an agent interacts with many agents, an agent can pursue only one negotiation at a time in some models. An agent has to terminate a current negotiation in disagreement first, and then pursue a more attractive outside alternative. This kind of model is called bilateral negotiation with *outside options* [87, 100]. However, the presumption that an agent can pursue only one negotiation at a time appears to be restrictive. In one-to-many negotiation [13, 14, 28, 96, 97, 98, 106, 128, 130], an agent can concurrently negotiate with multiple trading partners and an agent's proposal to one trading partner is affected by the status of its negotiation with other trading partners. In this work, each agent concurrently negotiates with multiple trading partners for multiple resources and an agent's proposals to each trading partner depends on the negotiation with all the trading partners.

Concurrent negotiations: Sim *et al.* [128, 130] proposed a coordination strategy for multi-resource negotiation where an agent can negotiate with multiple agents as in this work. Each buyer in [128, 130] knows the reserve price of each resource in advance and the buyer just needs to decide the concession strategy for each one-to-many negotiation for one resource. However, it is noted since [128, 130] focused on designing a concurrent mechanism for Grid resource co-allocation, the mechanism did not assume that consumer agents know the number of competing consumers. In contrast, each buyer in this work is assumed to only know the value of its high level task, i.e., the reserve price of all resources required for the high level task. We propose a set of heuristics for dynamically determining the reserve price of each resource based on the status of all negotiations. Furthermore, a buyer in [128, 130] only makes one

tentative agreement but in this work, a buyer may make more than one tentative agreement.

Organizational negotiation: Zhang *et al.* have studied a number of sophisticated negotiation problems in organizational contexts [137, 138]. Automated negotiation becomes increasingly complex and difficult as 1) agents are large-grained and complex with multiple goals and tasks, 2) agents often have more negotiation tasks and organizational relationships among heterogeneous agents become more complex, 3) negotiation process is tightly interleaved with agents' negotiation, scheduling and planning processes. Zhang *et al.* [137, 138] focus more on the coordination (a good "fit") of multiple negotiation tasks in organization context and they do not address agents' bargaining strategy in complex negotiation environments. In contrast, our work investigates how agents make concessions in dynamic negotiation environments where agents have multiple resources to negotiate. Zhang *et al.* [137] also considered multi-linked negotiation problems in which an agent needs to negotiate with multiple other agents about different subjects, and the negotiation over one subject has influence on negotiations over other subjects. They present a heuristic search algorithm for finding a near-optimal ordering of negotiation issues and their parameters.

Leveled commitment contracts: Sandholm *et al.* propose leveled-commitment contracts [121] in which the level of commitment is set by decommitting penalties. However, they only study the two-player game and they didn't investigate agents' bargaining strategies with decommitment from agreements. In addition, the problem setting in [121] is far from the real-world settings since they make strong assumptions about agents' knowledge such as outside options in the future. In the negotiation management system for CLASP [5], resource consumers can decommit from agreements made before at the cost of paying a penalty. However, the focus in their work is only on the scheduling problem. This work in contrast focuses on agents' negotiation strategies given that agents can decommit from agreements. Nguyen and Jennings [97, 98]

provide and evaluate a commitment model for concurrent negotiation. However, the maximum number of tentative agreements is determined prior to negotiation. In our work, the maximum number of tentative agreements is determined by the current market situation and will change dynamically during negotiation. In addition, our work studies a multi-resource negotiation problem, rather than single resource negotiation as in [97, 98]. Furthermore, Nguyen and Jennings [97, 98] make very restrictive assumptions about agents' available information, e.g., each agent is assumed to have knowledge about 1) other agents' negotiation strategies, 2) its negotiation success rate when it adopts certain strategy, and 3) its payoff when it adopts certain strategy. In this work, we assume that each agent has no knowledge about negotiation outcomes.

Learning in negotiation: A negotiating agent may have limited knowledge about others. Thus it is important for an agent to have the ability to update its beliefs based on its interactions with others. A variety of learning techniques have been used for building negotiation agents. Genetic algorithms (GAs) have been widely applied to automated negotiation. In general, GAs are used to enhance automated negotiation in two ways: 1) GAs were used as a decision making component at every round, e.g., [85], and 2) GAs were used to learn the best strategies, e.g., [13, 91]. Zeng and Sycara [136] present a sequential negotiation model and address multi-agent learning issues by explicitly modeling beliefs about the negotiation environment and the participating agents under a probabilistic framework using a Bayesian learning representation and updating mechanisms. Coehoorn and Jennings [37] showed that the preferences of a negotiation opponent in bilateral multi-issue negotiations can be effectively learnt by using kernel density estimation.

Mediation based negotiation: Another approach to resolve negotiation agents' conflicts is mediation. Ehtamo *et al.* [46] present a mediation-based negotiation framework for making trade-offs between cooperative negotiation agents. Klein *et al.* [73] presents a mediator based approach to negotiate complex contracts based on a ran-

dom searching method. The mediator in this model makes proposals to both agents. There are two types of negotiators: hill-climber and simulated annealer. Lai *et al.* [80] consider self-interested agents and propose a non-biased mediator who applies query learning to maintain near Pareto-efficiency without heavy computation. The major limitation of mediation based negotiation is that it ignores agents' strategic behavior of selfish negotiation agents. For instance, an agent may not truthfully report its information to the mediator in order to manipulate the mediation.

2.4 Auction Mechanism

Another class of widely used resource allocation mechanisms is the auction in which agents bid for the best resources. There are many possible auction designs depending on issues such as the efficiency of a given auction design, optimal and equilibrium bidding strategies, and revenue comparison. There are traditionally four types of auction that are used for single item allocation: 1) First-price sealed-bid auctions in which bidders simultaneously submit their bids in a sealed envelope. The individual with the highest bid wins and pays its bidding price. 2) Second-price sealed-bid auctions (or Vickrey auctions [134]) which is similar to first-price sealed-bid auctions except that the winner pays a price equal to the exact amount of the second highest bid. 3) English auctions in which the price is steadily raised by the auctioneer with bidders dropping out once the price becomes too high. This continues until there remains only one bidder who wins the auction at the current price. 4) Dutch auctions in which the auctioneer begins with a high asking price which is lowered until some participant is willing to accept the price. The winning participant pays the last announced price.

A combinatorial auction is an auction in which participants can place bids on combinations of discrete items, or "packages," rather than just individual items or continuous quantities (see the recent book [42] for more information about combi-

natorial auctions). Combinatorial auctions have been widely used in many applications such as sourcing in Internet commerce [114]. Combinatorial auctions present challenges compared to traditional auctions. In combinatorial auction schemes, a centralized controlling agent (the “auctioneer”) assumes responsibility for determining which agents receive which resources based on the bids submitted by individual agents. However, the problem of deciding successful bids, i.e., winner determination problem, is \mathcal{NP} -hard [110], meaning that a polynomial-time algorithm to find the optimal allocation is unlikely ever to be found. In addition, the auctioneer may face significant computational overload due to a large number of bids with complex structures.

Our negotiation approach for multi-agent resource allocation is of a *distributed* nature. In general, the allocation procedure used to find a suitable allocation of resources could be either *distributed* or *centralized*, e.g., combinatorial auctions. One of the most important arguments against centralized approaches is that it may be difficult to find an agent that could assume the role of an “auctioneer”. For instance, selfish resource providers may not trust the auctioneer and are not willing to comply with the decisions made by the auctioneer. In distributed approaches like automated negotiation, on the other hand, allocations emerge as the result of a sequence of distributed negotiations and each selfish agent acts on behalf of itself. The distributed model seems more natural in cases where resources belong to different selfish agents and finding optimal allocations may be (computationally) infeasible.

In dynamic resource allocation problems such as cloud computing, agents need to reason about future events while making decisions. For such dynamic resource allocation problem, strategy-proof mechanisms such as the well-known VCG auction is not necessarily strategy-proof and do not necessarily result in the most efficient usage [101]. While efficient online mechanisms have been proposed for dynamic environments, they only work in constrained settings and often rely on strong assumptions

about agents' knowledge (see [21, 101] for a survey). Against this background, we use distributed negotiation for dynamic resource allocation and compare it with some representative existing auction mechanisms.

2.5 Summary

In this chapter, we presented a brief overview of agent mediated negotiation as well as some market mechanisms. We first introduce the work on bargaining mechanism design and strategic bargaining games in economics. Then we discussed the related work on designing negotiation agents in AI. We also compared our work in this thesis and the state of the art. Analyzing agents' rational strategies for incomplete information bargaining and building agents for complex multi-resource negotiation are important for negotiation based resource allocation problems, but the literature does not provide satisfactory solutions. The focus of this thesis is on both challenging problems in bargaining theory and development of negotiation agents for practical complex resource allocation problems.

CHAPTER 3

NEGOTIATION WITH UNCERTAIN RESERVE PRICES

This chapter presents the algorithm for computing sequential equilibrium in pure strategies for bilateral bargaining with one-sided uncertainty regarding agents' reserve prices. There is no existing approach to solve the problem formally in the literature. Our approach is general in that it can be applied to other dynamic games with continuous strategy space.

3.1 Background

While there are many negotiation settings in electronic commerce transactions, the most common one (also the simplest one) is *bilateral* negotiation with a *single* negotiation issue. For instance, consider a scenario in which a buyer and a seller negotiate on the price of a good. In such a bargaining scenario, the two agents have different preferences over agreements. Thus agents need to make concessions toward a mutually acceptable agreement through a series of offers and counter offers. The negotiation fails if the two agents fail to make an agreement. There are many real-world negotiation examples such as the negotiation between a service provider and a customer over the price and the quality for providing a service.

In a bargaining game, an agent's strategy can be either pure or mixed. A pure strategy deterministically prescribes one action at any decision node of the game. In contrast, a mixed strategy is an assignment of a probability to each pure strategy. The concept of mixed strategies is very useful for games having no pure strategy equilibrium. However, the concept of mixed strategies has been criticized for being

“intuitively problematic” since randomization lacks behavioral support [17]. When mixed strategies are considered, the number of sequential equilibria of the game usually increases and coordination problems of choosing a equilibrium strategy profile has not been fully addressed. Due to these reasons, we focus on pure strategies equilibrium. Fortunately, simulation results show that there is at least one pure strategy sequential equilibrium in 99.7% of various bilateral bargaining scenarios we will look at in our experiments. Additionally, experimentally we found that as the number of uncertain types and deadlines increase, all cases had at least one sequential equilibrium.

One major motivation of the study of negotiation theory is designing successful negotiation agents in practical applications. However, it is often impossible to compute agents’ rational strategies for more realistic complex negotiation games where there are many agents and a large number of uncertainties. While we study a relatively simple negotiation problem in this chapter, our analysis can give us guidelines for designing negotiation agents for practical negotiation problems. In addition, as will be discussed at the end of this chapter, there may be ways to extend the solution approaches presented in this chapter to compute agents’ equilibrium strategies in more complex negotiation games.

The rest of this chapter is organized as follows. We start with complete information negotiation in Section 3.2. Section 3.3 discusses the difficulty caused by introducing uncertainty. Section 3.4 introduces our algorithm. Section 3.5 shows how to compute the buyer’s equilibrium offer and Section 3.6 shows how to compute the seller’s equilibrium offer. Section 3.7 analyzes equilibrium existence. Section 3.8 compared agents’ utilities while using different strategies. Section 3.9 discusses two potential applications of our approach. Section 3.10 concludes this chapter and outlines future research directions.

3.2 Bargaining with Complete Information

This section describes the discrete time bargaining between a buyer \mathbf{b} and a seller \mathbf{s} . The seller wants to sell a single indivisible good to the buyer with a price. All the agents enter the market at time 0. An alternating-offers bargaining protocol is utilized. Formally, the buyer \mathbf{b} and the seller \mathbf{s} can act at times $t \in \mathbb{N}$. The player function $\iota : \mathbb{N} \rightarrow \{\mathbf{b}, \mathbf{s}\}$ returns the agent that acts at time t and is such that $\iota(t) \neq \iota(t+1)$, i.e., a pair of agents bargain by making offers in alternate fashion. This chapter focuses on single-issue negotiation but this model can be easily extended to handle multi-issue negotiation [61].

Possible actions $\sigma_{\iota(t)}^t$ of agent $\iota(t)$ at any time point $t > 0$ are:

1. *offer*[x], where $x \in \mathbb{R}$ is the proposed price for the good;
2. *exit*, which indicates that negotiation fails;
3. *accept*, which indicates that \mathbf{b} and \mathbf{s} have reached an agreement.

At time point $t = 0$ the only allowed actions are 1) and 2). If $\sigma_{\iota(t)}^t = \textit{accept}$ the bargaining stops and the outcome is $o = (x, t)$, where x is the value such that $\sigma_{\iota(t-1)}^{t-1} = \textit{offer}[x]$. This is to say that the agents agree on the value x at time point t . If $\sigma_{\iota(t)}^t = \textit{exit}$ the bargaining stops and the outcome is *FAIL*. Otherwise the bargaining continues to the next time point.

Each agent $\mathbf{a} \in \{\mathbf{b}, \mathbf{s}\}$ has a utility function $U_{\mathbf{a}} : (\mathbb{R} \times \mathbb{N}) \cup \textit{FAIL} \rightarrow \mathbb{R}$, which represents its gain over the possible bargaining outcomes. Each utility function $U_{\mathbf{a}}$ depends on \mathbf{a} 's reserve price $\text{RP}_{\mathbf{a}} \in \mathbb{R}^+$, temporal discount factor $\delta_{\mathbf{a}} \in (0, 1)$,¹ and deadline $T_{\mathbf{a}} \in \mathbb{N}, T_{\mathbf{a}} > 0$. If the bargaining outcome is $o = (x, t)$, then the utility function $U_{\mathbf{a}}$ is defined as:

¹A discount factor is used to model bargaining cost, which is a common assumption in the bargaining literature [111, 112, 51, 61].

$$U_{\mathbf{a}}(x, t) = \begin{cases} (\text{RP}_{\mathbf{a}} - x) \cdot (\delta_{\mathbf{a}})^t & \text{if } t \leq T_{\mathbf{a}} \text{ and } \mathbf{a} \text{ is a buyer} \\ (x - \text{RP}_{\mathbf{a}}) \cdot (\delta_{\mathbf{a}})^t & \text{if } t \leq T_{\mathbf{a}} \text{ and } \mathbf{a} \text{ is a seller} \\ \epsilon < 0 & \text{otherwise} \end{cases}$$

If the outcome is *FAIL*, $U_{\mathbf{a}}(\text{FAIL}) = 0$. Notice that the assignment of a strictly negative value to $U_{\mathbf{a}}$ after \mathbf{a} 's deadline allows one to capture the essence of the deadline: an agent, after its deadline, strictly prefers to exit the negotiation rather than to reach any agreement. Therefore, the bargaining model we consider is a finite horizon game. Finally, we assume the feasibility of the problem, i.e., $\text{RP}_{\mathbf{b}} \geq \text{RP}_{\mathbf{s}}$.

With complete information the appropriate solution concept for the game is the *subgame perfect equilibrium* in which agents' strategies are in equilibrium in every possible subgame [58]. Note that there is no deadline constraint in the negotiation protocol, which indicates that agents are allowed to offer and counteroffer also after their deadlines have expired. However, the deadline constraint is in both agents' utility functions such that no rational agent will continue negotiation after its deadline. Therefore, the bargaining game is a finite horizon game and the subgame perfect equilibrium can be found employing the backward induction method.

Initially, it is determined that the game rationally stops at time point $T = \min(T_{\mathbf{b}}, T_{\mathbf{s}})$. The equilibrium outcome of every subgame starting from $t \geq T$ is *FAIL*, since at least one agent will exit from bargaining. Therefore, at $t = T$ agent $\iota(T)$ would accept any offer x which gives it a utility not worse than *FAIL*, namely, any offer x such that $U_{\iota(T)}(x, T) \geq 0$. From $t = T - 1$ back to $t = 0$ it is possible to find the optimal offer agent $\iota(t)$ can make at t , if it makes an offer, and the offers that it would accept. $x^*(t)$ denotes the optimal offer of agent $\iota(t)$ at t . $x^*(t)$ is the offer such that, if $t < T - 1$, agent $\iota(t + 1)$ is indifferent at $t + 1$ between accepting it and rejecting it to make its optimal offer $x^*(t + 1)$ and, if $t = T - 1$, agent $\iota(t + 1)$ is indifferent at $t + 1$ between accepting it and exiting. Formally, $x^*(t)$ is such that

$U_{\iota(t+1)}(x^*(t), t) = U_{\iota(t+1)}(x^*(t+1), t+1)$ if $t < T - 1$ and $U_{\iota(t+1)}(x^*(t), t) = 0$ if $t = T - 1$. The offers agent $\iota(t)$ would accept at t are all those offers that give it a utility no worse than the utility given by offering $x^*(t)$. The equilibrium strategy of any subgame starting from $0 \leq t < T$ prescribes that agent $\iota(t)$ offers $x^*(t)$ at t and agent $\iota(t+1)$ accepts it at $t+1$.

Backward propagation is used to provide a recursive formula for $x^*(t)$: given value x and agent \mathbf{a} , we call backward propagation of value x for agent \mathbf{a} the value y such that $U_{\mathbf{a}}(y, t-1) = U_{\mathbf{a}}(x, t)$; we employ the arrow notation $x_{\leftarrow \mathbf{a}}$ for backward propagations. Formally, $x_{\leftarrow \mathbf{b}} = \text{RP}_{\mathbf{b}} - (\text{RP}_{\mathbf{b}} - x) \cdot \delta_{\mathbf{b}}$ and $x_{\leftarrow \mathbf{s}} = \text{RP}_{\mathbf{s}} + (x - \text{RP}_{\mathbf{s}}) \cdot \delta_{\mathbf{s}}$. If a value x is backward propagated n times for agent \mathbf{a} , we write $x_{\leftarrow n[\mathbf{a}]}$, e.g., $x_{\leftarrow 2[\mathbf{a}]} = (x_{\leftarrow \mathbf{a}})_{\leftarrow \mathbf{a}}$. If a value is backward propagated for more than one agent, we list them left to right in the subscript, e.g., $x_{\leftarrow \mathbf{b}2[\mathbf{s}]} = ((x_{\leftarrow \mathbf{b}})_{\leftarrow \mathbf{s}})_{\leftarrow \mathbf{s}}$. The values of $x^*(t)$ can be calculated recursively from $t = T - 1$ back to $t = 0$ as follows:

$$x^*(t) = \begin{cases} \text{RP}_{\iota(t+1)} & \text{if } t = T - 1 \\ (x^*(t+1))_{\leftarrow \iota(t+1)} & \text{if } t < T - 1 \end{cases}$$

It can be observed that $x_{\leftarrow \mathbf{b}} \geq x$ as $x_{\leftarrow \mathbf{b}} - x = \text{RP}_{\mathbf{b}} - (\text{RP}_{\mathbf{b}} - x) \cdot \delta_{\mathbf{b}} - x = (1 - \delta_{\mathbf{b}})(\text{RP}_{\mathbf{b}} - x) \geq 0$, and $x_{\leftarrow \mathbf{s}} \leq x$ as $x_{\leftarrow \mathbf{s}} - x = \text{RP}_{\mathbf{s}} + (x - \text{RP}_{\mathbf{s}}) \cdot \delta_{\mathbf{s}} - x = (\delta_{\mathbf{s}} - 1)(x - \text{RP}_{\mathbf{s}}) \leq 0$.

Figure 3.1 shows an example of backward induction construction with parameters $\text{RP}_{\mathbf{b}} = 100$, $\text{RP}_{\mathbf{s}} = 0$, $\iota(0) = \mathbf{s}$, $\delta_{\mathbf{b}} = 0.75$, $\delta_{\mathbf{s}} = 0.8$, $T_{\mathbf{b}} = 10$, and $T_{\mathbf{s}} = 11$. The backward induction process starts from time $T = \min\{T_{\mathbf{b}}, T_{\mathbf{s}}\} = 10$. At time 10, the seller is willing to accept any offer which is no less than its reserve price and thus the optimal offer at time $t = 9$ is $x^*(9) = \text{RP}_{\mathbf{s}} = 0$. The optimal offer of the seller at time $t = 8$ is $x^*(8) = (\text{RP}_{\mathbf{s}})_{\leftarrow \mathbf{b}} = \text{RP}_{\mathbf{b}} - (\text{RP}_{\mathbf{b}} - \text{RP}_{\mathbf{s}}) \cdot \delta_{\mathbf{b}} = 25$. Analogously, the optimal offer of the buyer at time $t = 7$ is $x^*(7) = (x^*(8))_{\leftarrow \mathbf{s}} = \text{RP}_{\mathbf{s}} + (x^*(8) - \text{RP}_{\mathbf{s}}) \cdot \delta_{\mathbf{s}} = 20$.




Figure 3.1. Backward induction construction with $RP_b = 100$, $RP_s = 0$, $\iota(0) = s$, $\delta_b = 0.75$, $\delta_s = 0.8$, $T_b = 10$, $T_s = 11$; at each time point t the optimal offer $x^*(t)$ is marked; the dashed lines are isoutility curves.

Following this procedure, we can get agents' optimal offers from time $t = 6$ to the initial time point $t = 0$.

Finally, agents' equilibrium strategies can be defined on the basis of $x^*(t)$ as follows:

$$\sigma_{\mathbf{b}}^*(t) = \begin{cases} t = 0 & offer[x^*(0)] \\ 0 < t < T & \begin{cases} \text{if } \sigma_{\mathbf{s}}(t-1) = offer[x] \text{ with } x \leq (x^*(t))_{\leftarrow \mathbf{b}} & \textit{accept} \\ \text{otherwise} & offer[x^*(t)] \end{cases} \\ T \leq t \leq T_{\mathbf{b}} & \begin{cases} \text{if } \sigma_{\mathbf{s}}(t-1) = offer[x] \text{ with } x \leq RP_{\mathbf{b}} & \textit{accept} \\ \text{otherwise} & \textit{exit} \end{cases} \\ T_{\mathbf{b}} < t & \textit{exit} \end{cases}$$

$$\sigma_{\mathbf{s}}^*(t) = \begin{cases} t = 0 & offer[x^*(0)] \\ 0 < t < T & \begin{cases} \text{if } \sigma_{\mathbf{b}}(t-1) = offer[x] \text{ with } x \geq (x^*(t))_{\leftarrow \mathbf{s}} & \textit{accept} \\ \text{otherwise} & offer[x^*(t)] \end{cases} \\ T \leq t \leq T_{\mathbf{s}} & \begin{cases} \text{if } \sigma_{\mathbf{b}}(t-1) = offer[x] \text{ with } x \geq RP_{\mathbf{s}} & \textit{accept} \\ \text{otherwise} & \textit{exit} \end{cases} \\ T_{\mathbf{s}} < t & \textit{exit} \end{cases}$$

We can see that the above strategies constitute a unique subgame perfect equilibrium of bargaining with complete information. The equilibrium can be found in time linear to the maximum deadline of the two agents. At the equilibrium, the two agents reach an agreement at time $t = 1$ and the agreement price is $x^*(0)$.

3.3 One-sided Uncertainty about Reserve Prices

In this section, we first loosen the complete information bargaining model in the previous section by introducing one-sided uncertainty regarding the buyer's reserve price. We then review the existing approaches in the literature.

3.3.1 Introducing Uncertainty

With uncertain information, the appropriate solution concept for an extensive-form game is *sequential equilibrium* [58]. A sequential equilibrium is a pair $a = \langle \mu, \sigma \rangle$ (also called an assessment) where μ is a belief system that specifies how agents' beliefs evolve during the game and σ specifies agents' strategies. At an equilibrium μ must be *consistent* with respect to σ and σ must be *sequentially rational* given μ . Informally, the rationality requirement says that after every possible sequence of actions, an agent's strategy must maximize its expected utility given its beliefs and its opponent's equilibrium strategy. An assessment a is consistent (in the sense of Kreps and Wilson [79]) if there exists a sequence of totally mixed strategy profiles (with associated sensible beliefs updated according to Bayes' rule) that converges to the equilibrium profile.

We assume the one-sided uncertainty regarding the type of the buyer \mathbf{b} (the case of having uncertainty with the type of the seller \mathbf{s} can be analyzed analogously). The buyer \mathbf{b} can be of finitely many types $\{\mathbf{b}_1, \dots, \mathbf{b}_n\}$ in which buyer type \mathbf{b}_i has an associated reserve price RP_i . The initial belief of \mathbf{s} on \mathbf{b} is $\mu(0) = \langle \Delta_{\mathbf{b}}^0, P_{\mathbf{b}}^0 \rangle$ where $\Delta_{\mathbf{b}}^0 = \{\mathbf{b}_1, \dots, \mathbf{b}_n\}$ and $P_{\mathbf{b}}^0 = \{\omega_{\mathbf{b}_1}^0, \dots, \omega_{\mathbf{b}_n}^0\}$ such that $\sum_i \omega_{\mathbf{b}_i}^0 = 1$. $\omega_{\mathbf{b}_i}^0$ is the *priori* probability that \mathbf{b} is of type \mathbf{b}_i . The belief of \mathbf{s} on the type of \mathbf{b} at time t is $\mu(t)$. The probability assigned by \mathbf{s} to $\mathbf{b} = \mathbf{b}_i$ at time t is denoted $\omega_{\mathbf{b}_i}^t$. Given an assessment $a = \langle \mu, \sigma \rangle$, there are multiple possible bargaining outcomes: outcome $o_{\mathbf{b}_i}$ if $\mathbf{b} = \mathbf{b}_i$. We denote bargaining outcome as $o = \langle o_{\mathbf{b}_1}, \dots, o_{\mathbf{b}_n} \rangle$.

Seller \mathbf{s} 's belief of the type of buyer \mathbf{b} will evolve based on its observed actions and the buyer's equilibrium strategies. On the equilibrium path, \mathbf{s} 's belief at any time t is $\mu(t) = \langle \Delta_{\mathbf{b}}^t, P_{\mathbf{b}}^t \rangle$. As is customary in economic studies [112], we consider only stationary systems of beliefs, i.e., if \mathbf{s} believes a \mathbf{b} 's type with zero probability at time point t , it will continue to believe such a type with zero probability at any time point $t' > t$. We can therefore specify $\mu(t)$ by specifying $\Delta_{\mathbf{b}}^t$. Moreover, given

that $\mu(t) = \Delta_{\mathbf{b}}^t$ and we only consider pure strategies, the probability that \mathbf{b} is of type $\mathbf{b}_i \in \Delta_{\mathbf{b}}^t$ is $\omega_{\mathbf{b}_i}(\Delta_{\mathbf{b}}^t) = \frac{\omega_{\mathbf{b}_i}^0}{\sum_{\mathbf{b}_j \in \Delta_{\mathbf{b}}^t} \omega_{\mathbf{b}_j}^0}$.

We need to also specify the belief system off the equilibrium path, i.e., when an agent takes an action that is not optimal. We use the *optimistic conjectures* [112].² That is, when buyer \mathbf{b} acts off the equilibrium strategy, agent \mathbf{s} will believe that agent \mathbf{b} is of its “weakest” type, i.e., the type against which the seller would gain the most. This choice is made to assure the existence of the equilibrium for the largest subset of the space of the parameters [61]. In our case, the weakest type is the buyer type with the highest reserve price (see Section 3.4.4 for the proof). That is, if $\mu(t-1) = \Delta_{\mathbf{b}}^{t-1}$ and \mathbf{b} acts off the equilibrium strategy at time $t-1$, it follows that $\Delta_{\mathbf{b}}^t = \mathbf{b}_h(\Delta_{\mathbf{b}}^{t-1})$ where $\mathbf{b}_h(\Delta_{\mathbf{b}}^{t-1})$ is the buyer type with the highest reserve price in buyer types $\Delta_{\mathbf{b}}^{t-1}$.

3.3.2 Existing Solutions in Literature

Computation agents’ equilibrium strategies of an extensive-form game with imperfect information is well known to be hard and classic game theory does not provide any general approaches to find sequential equilibria. While there has been long standing literature in solving bargaining games with uncertainty since Rubinstein’s path-breaking work [112], there is no existing approach that can be applied to solve the bargaining problem studied in this chapter. An interested reader can find a more detailed survey on bargaining with uncertainty in [18].

Computer science researchers have proposed a number of algorithms for computing Nash equilibria (e.g., [62, 63, 74]) or sequential equilibria (e.g., [92]). However, these algorithms are not applicable in solving bargaining games since they only consider finite strategy space rather than continuous strategy space (i.e., price) considered in

²While this chapter assumes optimistic conjectures, our approach can be used for any belief update rules for agents’ actions off the equilibrium path.

this chapter. Due to the same reason, enumeration based methods (e.g., [104]) cannot be applied to our bargaining problem as well.

The microeconomic literature provides a number of results for some specific bargaining problems with uncertainty. For instance, Rubinstein [112] considered bilateral infinite horizon bargaining with uncertainty over two possible discount factors. With the unrealistic infinite horizon assumption, Rubinstein found a number of closed form results such as how the discount factors will affect the equilibrium outcome. Sandholm and Vulkan [117] analyze agents' strategic behavior in a slight variation of the war-of-attrition game where the surplus can be divided. They consider a finite horizon alternating-offer bilateral bargaining game where agents have uncertain deadlines, time is continuous, and there are not discount factors. In contrast, Gatti *et al.* [61] relaxed the infinite horizon deadlines and provided an algorithm to compute agents' equilibrium strategies in bilateral bargaining with one-sided uncertain deadlines. They proved that agent types would adopt the same strategy at any time point before their deadlines, which may not be true in our case with uncertain reserve prices. Therefore, their approach cannot be applied to our case. Cramton [39] considered a special infinite horizon bargaining protocol in which only the seller can make offers and the buyer can only accept or reject the seller's offer. Chatterjee and Samuelson [34, 35] studied bilateral infinite horizon bargaining with two-type uncertainty over the reservation values. The absence of agents' deadlines makes these two results nonapplicable to the situation we study in the chapter. An *et al.* [6] only considered two-type uncertainty about reserve prices and their approach cannot be directly extended to handle multiple types. The presence of multi-type uncertainty increases the computational complexity of the procedure to find equilibrium strategies and requires more stringent equilibrium existence conditions.

The only known general approach that be potentially applied to our bargaining problem is the backward induction approach by Fatima *et al.* [51, 52]. They studied




Figure 3.2. Failure of the approach in [51, 52] with $T = 5$, $\iota(0) = s$, $RP_s = 10$, $RP_1 = 90$, $RP_2 = 70$, $\omega_{b_1}^0 = 0.8$, $\omega_{b_2}^0 = 0.2$, $\delta_s = 0.7$, and $\delta_b = 0.8$; agents' offers in complete information settings were also showed.

the alternating-offers protocol with multiple negotiation issues and uncertainty over the weights of the issues. They proposed an algorithm based on backward induction to compute sequential equilibria. Note that as in this chapter, Fatima *et al.* [51, 52] also focus on pure strategy equilibria. Basically, their algorithm searches in the space of the strategies exploiting the backward induction from the last possible deadline to $t = 0$ with agents' initial beliefs, and, once the optimal strategies at time point $t = 0$ have been found, the system of beliefs is designed to be consistent with them. It has been shown in [61] through a counter example that unfortunately an equilibrium return by their algorithm is not necessary a sequential equilibrium in pure strategies for bilateral bargaining with uncertain deadlines.

We show a simple example where the algorithm in [52, 51] fails in the bargaining problem studied in this chapter (see Figure 3.2 for agents' equilibrium offers computed by their algorithm and agents' equilibrium offers in complete information settings). Consider the following scenario: $T = 5$, $\iota(0) = \mathbf{s}$, $RP_{\mathbf{s}} = 10$, $RP_1 = 90$, $RP_2 = 70$, $\omega_{\mathbf{b}_1}^0 = 0.8$, $\omega_{\mathbf{b}_2}^0 = 0.2$, $\delta_{\mathbf{s}} = 0.7$, and $\delta_{\mathbf{b}} = 0.8$. Let $x_{\mathbf{b}_i}^*(t)$ be any agent optimal offer at time t when buyer \mathbf{b} is of type \mathbf{b}_i in the complete information setting. Let $x^f(t)$ be any agent optimal offer at time t computed by the algorithm in [52, 51]. Agents' equilibrium offers are computed with the initial belief. At time $t = 4$, seller \mathbf{s} can offer either $x_{\mathbf{b}_1}^*(4) = RP_1 = 90$ which gives the seller a utility of $0.8 \cdot (90 - 10) \cdot 0.7^5 = 10.75648$ or $x_{\mathbf{b}_2}^*(4) = RP_2 = 70$ which give the seller a utility of $(70 - 10) \cdot 0.7^5 = 10.0842$. Therefore, the seller's optimal offer at time $t = 4$ is $x^f(4) = 90$ and the equivalent price is 74. Then both buyer types' optimal offer at time $t = 3$ is $x^f(3) = (74)_{\leftarrow \mathbf{s}} = 54.8$. At time $t = 2$, seller \mathbf{s} can offer either $(54.8)_{\leftarrow \mathbf{b}_1} = 61.84$ or $(54.8)_{\leftarrow \mathbf{b}_2} = 57.84$. According to [52, 51], offering price 57.84 will be accepted by both buyer types and thus the seller can gain a utility of $(57.84 - 10) \cdot 0.7^3 = 16.40912$. In contrast, offering price 61.84 will only be accepted by buyer type \mathbf{b}_1 and the seller's equilibrium offer at time $t = 4$ will be accepted. Thus, offering price 61.84 will the seller a utility of $0.8 \cdot (61.84 - 10) \cdot 0.7^3 + 0.2 \cdot (74 - 10) \cdot 0.7^5 = 16.376192$. Thus the optimal offer of the seller at time $t = 2$ is $x^f(2) = 57.84$. Then both buyer types' optimal offer at time $t = 1$ is $x^f(1) = (57.84)_{\leftarrow \mathbf{s}} = 43.488$. The seller at time $t = 0$ can offer either $(43.488)_{\leftarrow \mathbf{b}_1} = 52.7904$ or $(43.488)_{\leftarrow \mathbf{b}_2} = 48.7904$ and its optimal offer is $x^f(0) = 52.7904$ which will only be accepted by buyer type \mathbf{b}_1 . According to [52, 51], buyer type \mathbf{b}_1 will accept the offer 52.7904 at time $t = 0$ since the optimal offer 52.7904 is \mathbf{b}_1 's backward propagated value of its 43.488 at time $t = 1$. Accordingly, the seller will update its belief as follows: if its optimal offer 52.7904 is rejected, it will update its belief to $\{\mathbf{b}_2\}$.

However, the above strategy profile is not in sequential equilibrium since buyer \mathbf{b}_1 has an incentive to reject the seller’s equilibrium offer at time $t = 1$ (also see Figure 3.2). If buyer type \mathbf{b}_1 rejects the offer 52.7904 at time $t = 1$ and makes a counter offer 41.92, the seller will accept it since 41.92 is buyer’s equilibrium offer when the buyer is of type \mathbf{b}_2 . By doing so, buyer \mathbf{b}_1 gains a utility of $(90 - 41.92) \cdot 0.8^2 = 30.7712$ which is higher than its utility $(90 - 52.7904) \cdot 0.8 = 29.76768$ when it accepts the seller’s equilibrium offer 52.7904. As pointed out in [61], the reason behind the failure of [52, 51] in producing equilibrium strategies for some settings of parameters is that in each step of backward induction they limit the search to the space of the strategies, but they do not verify the existence of a consistent system of beliefs such that the found strategy is sequentially rational. In other words, they break the circularity of strategies and belief systems. In the above example, they decide the acceptance price of buyer type \mathbf{b}_1 with the initial belief and ignore the effect of the seller’s belief update rule. As a result, once their algorithm has produced the agents’ strategies at $t = 0$ and has designed the system of beliefs consistent with them, the strategies may not be sequentially rational given the designed system of beliefs.

3.4 The Algorithm for Finding All Sequential Equilibria

This section first introduces the high level idea of our approach. Following that we analyze some observations that can be used to drastically reduce the required computation based on our basic approach. Finally we introduce the algorithm for finding all sequential equilibria of a bilateral bargaining game with one-sided uncertainty.

3.4.1 High Level Idea of the Approach

Our approach follows the spirit of *backward induction*: To compute agent \mathbf{a} ’s equilibrium offer with belief $\Delta_{\mathbf{b}}$ at time $t < T - 1$, agent \mathbf{a} takes into account all the sequential equilibria in the *continuation game* with different beliefs starting from time




Figure 3.3. A high level illustration of our approach ($\iota(t) = \mathbf{s}$ and $|\Delta^0| > 1$)

$t + 1$. A continuation game is composed of an information set for one agent (buyer or seller) and all of its successor nodes from the original bilateral bargaining game. Note that there is no subgame for the bargaining game with uncertainty. There are continuation games starting from time points $0, 1, \dots$. Let $\Gamma(t)$ be the continuation game starting from time t . In the continuation game $\Gamma(t)$, agent $\iota(t)$ makes its offer at time t first. Let $\Gamma(t, \Delta_{\mathbf{b}})$ be the continuation game $\Gamma(t)$ with seller \mathbf{s} 's initial belief $\Delta_{\mathbf{b}}$. The problem of finding sequential equilibria for a bargaining problem is finding all sequential equilibria for the continuation game $\Gamma(0, \Delta_{\mathbf{b}}^0)$.

The definition of a sequential equilibrium requires that after observing buyer \mathbf{b} 's counter offer at time t , seller \mathbf{s} must update its belief about \mathbf{b} 's type using a *belief update rule*. The counter offer of buyer \mathbf{b} at time t indicates buyer \mathbf{b} 's following actions: 1) seller \mathbf{s} 's last offer is rejected by the buyer \mathbf{b} if $t > 0$, and 2) buyer \mathbf{b} makes a new offer at time t . Seller \mathbf{s} will update its belief given all the actions of

buyer **b**. Therefore, there are two types of belief update rules: 1) *reject update rules* applied when buyer **b** rejects seller **s**'s offer and *offer update rules* applied when buyer **b** makes a new offer. Assume seller **s**'s belief before proposing its offer x is $\Delta_{\mathbf{b}}$, the reject update rule is of the following form: If x is rejected, **s**'s belief about the type of buyer **b** is $\Delta'_{\mathbf{b}} \subseteq \Delta_{\mathbf{b}}$. Similarly, the offer update rule has the following form: If buyer **b** offers x , seller **s**'s belief about the type of **b** is $\Delta'_{\mathbf{b}} \subseteq \Delta_{\mathbf{b}}$ where $\Delta_{\mathbf{b}}$ is **s**'s belief before applying the offer update rule. After receiving buyer **b**'s offer at time $t = 0$, **s** will only apply the offer update rule. In any other situation (i.e., buyer **b** first rejects **s**'s offer and then makes a new offer), seller **s** will apply the reject update rule first and then apply the offer update rule.³

While the seller is making an offer at time t given the sequential equilibria for the continuation game $\Gamma(t + 1)$ with different beliefs, the seller will consider different reject update rules and compute its equilibrium offer for each rule. With pure strategies, the seller's reject update rules are finite. The other situation is deciding the buyer's equilibrium offer at time t given the sequential equilibria for the continuation game $\Gamma(t + 1)$ with different beliefs, the buyer will consider different *choice rules* regarding whether different buyer types behave in the same way or behave in different ways. With pure strategies, buyer types' choice rules are finite. For each choice rule, we compute each buyer type's optimal offer and its corresponding offer update rule. While computing agents' equilibrium strategies, we also construct equilibrium existence conditions and check whether those conditions are satisfied.

Roughly, the idea of our approach is the following (see Figure 3.3). To compute agents' equilibrium offers at a continuation game, we first compute sequential equilibria in its continuation game with different beliefs. Then we compute agents'

³This belief update process is obvious when an agent is required to send a rejection message before making a counter offer [51]. For the sake of simplicity, a rejecting agent does not need to send a rejection message in our protocol.

equilibrium offers together with agents' belief update rules. There are two cases. While computing the seller's equilibrium strategy, we enumerate all possible reject update rules (e.g., reject update rules 1 and 2 in Figure 3.3) and for each reject update rule, we first compute the seller's optimal strategy in the corresponding continuation game. For example, for the reject update rule 1 in Figure 3.3, we first solve the continuation game $\Gamma(t + 1, \Delta^1)$ where $\Delta^1 \subseteq \Delta^0$ is the seller's updated belief if the seller's offer is rejected. While computing the buyer's equilibrium strategy, we consider all choice rules and compute different buyer types' optimal offer for each choice rule. For instance, for the choice rule 3, we need to first solve the continuation game $\Gamma(t + 2, \Delta^3)$. There are two processes involved in computing all sequential equilibria: a *forward search* process to determine the set of continuation games to solve and a *backward induction* process to compute agents' equilibrium strategies based on all sequential equilibria of continuation games. Furthermore, we introduce some equilibrium existence conditions: if they are satisfied, there is a sequential equilibrium in the continuation game.

Take the bargaining problem in Figure 3.2 as an example. Our objective is to compute all sequential equilibria for the continuation game $\Gamma(0, \{\mathbf{b}_1, \mathbf{b}_2\})$. Since $\iota(0) = \mathbf{s}$, we need to consider different reject update rules. Consider the reject update rule that the seller is making an offer x that will only be accepted by buyer type \mathbf{b}_1 , i.e., if the buyer rejects offer x , the seller will update its belief to $\{\mathbf{b}_2\}$. To compute the optimal offer x at time $t = 0$, we first compute all sequential equilibria for the continuation game $\Gamma(1, \{\mathbf{b}_2\})$ starting from time $t = 1$. For another reject update rule that the seller is making an offer x that will be rejected by both buyer types, we need to first compute sequential equilibria for the continuation game $\Gamma(1, \{\mathbf{b}_1, \mathbf{b}_2\})$ with the original belief. To compute sequential equilibria for the continuation game $\Gamma(1, \{\mathbf{b}_1, \mathbf{b}_2\})$, we need to consider buyer types' different choice rules. Consider the choice rule that buyer type \mathbf{b}_1 makes an acceptable offer but buyer type \mathbf{b}_2 makes an

offer that will be rejected. For this choice rule, we need to first compute sequential equilibria for continuation games $\Gamma(2, \{\mathbf{b}_1\})$ and $\Gamma(2, \{\mathbf{b}_2\})$ starting from time $t = 2$. In the same way, we can recursively try different choice rules and reject update rules to compute all sequential equilibria of the bargaining game.

3.4.2 Computation Reduction

This section provides some theoretical results which drastically reduce the computation complexity. In an equilibrium, it is possible that the seller will make an offer that will be rejected by all the buyer types. Without loss of generality, we assume ϖ be seller's offer that will be rejected by all buyer types. Assume that the seller's belief is $\Delta_{\mathbf{b}}$. A reject update rule specifies the seller's updated belief $\Delta'_{\mathbf{b}} \subseteq \Delta_{\mathbf{b}}$ if the seller's offer is rejected. Therefore, the number of reject update rules are finite since the number of belief set $\Delta'_{\mathbf{b}} \subseteq \Delta_{\mathbf{b}}$ is no more than $2^{|\Delta_{\mathbf{b}}|}$. However, there is no sequential equilibrium for most of the reject update rules.

Theorem 1. *If there is a reject update rule with updated belief $\Delta'_{\mathbf{b}} \subseteq \Delta_{\mathbf{b}}$ such that $RP_i < RP_j$ for buyer type $\mathbf{b}_i \in \Delta_{\mathbf{b}} \setminus \Delta'_{\mathbf{b}}$ and buyer type $\mathbf{b}_j \in \Delta'_{\mathbf{b}}$, agents' strategies are not sequentially rational.*

Proof. This result can be proved by contradiction. If there is a sequential equilibrium with this reject update rule in which \mathbf{s} 's equilibrium offer at time t is x , the following two conditions are satisfied: 1) \mathbf{b}_i has no incentive to behave as \mathbf{b}_j , i.e., $U_{\mathbf{b}_i}(x, t+1) \geq U_{\mathbf{b}_i}(e_{\mathbf{b}_j}^{t+1}|\Delta'_{\mathbf{b}}, t+2)$ where $e_{\mathbf{b}_j}^{t+1}|\Delta'_{\mathbf{b}}$ is \mathbf{b}_j 's equivalent offer (will be defined later) in the continuation game starting from $t+1$ with belief $\Delta'_{\mathbf{b}}$; 2) \mathbf{b}_j has no incentive to behave as \mathbf{b}_i , i.e., $U_{\mathbf{b}_j}(e_{\mathbf{b}_j}^{t+1}|\Delta'_{\mathbf{b}}, t+2) \geq U_{\mathbf{b}_j}(x, t+1)$. Condition 1) suggests that $x \leq (e_{\mathbf{b}_j}^{t+1}|\Delta'_{\mathbf{b}})_{\leftarrow \mathbf{b}_i}$ and condition 2) indicates that $x \geq (e_{\mathbf{b}_j}^{t+1}|\Delta'_{\mathbf{b}})_{\leftarrow \mathbf{b}_j}$. Therefore, equilibrium existence conditions requires that $(e_{\mathbf{b}_j}^{t+1}|\Delta'_{\mathbf{b}})_{\leftarrow \mathbf{b}_j} \leq (e_{\mathbf{b}_j}^{t+1}|\Delta'_{\mathbf{b}})_{\leftarrow \mathbf{b}_i}$, which cannot be true since $RP_i < RP_j$. \square

Due to Theorem 1, we only need to consider reject update rules in which buyer types with higher reserve prices accept the seller's equilibrium offer while buyer types with lower reserve prices reject the seller's equilibrium offer. Assume that the seller's belief at a time point is $\Delta_{\mathbf{b}}$. The total number of reject update rules we need to consider is at most $|\Delta_{\mathbf{b}}|$ rather than $2^{|\Delta_{\mathbf{b}}|}$. For each reject update rule at time t , we need to first compute the sequential equilibrium for the continuation game $\Gamma(t+1)$ with the corresponding *reasonable* updated belief if the seller's offer is rejected, i.e., $\Delta'_{\mathbf{b}}$. Accordingly, we need to compute sequential equilibria for the continuation game with at most $|\Delta_{\mathbf{b}}|$ different reasonable beliefs.

In addition to the above rejected update rules in which according to the equilibrium strategy at least one buyer type will reject the seller's offer, we also need to consider the case that according to the equilibrium strategy, the seller's offer will be accepted by all buyer types. If the offer is rejected (i.e., the buyer is acting off the equilibrium path), the seller will update its belief to the buyer type with the highest reserve price according to the optimistic conjectures. We call this reject update rule as *null* reject update rule.

The other situation is deciding the buyer's equilibrium offer at time t . We use the term "choice rule" to characterize buyer types' strategies regarding whether they behave in the same way at a specific decision making point. With pure strategies, buyer types' choice rules are finite. Consider that the belief of \mathbf{s} on the type of \mathbf{b} at time t is $\mu(t) = \Delta_{\mathbf{b}}$ where $|\Delta_{\mathbf{b}}| > 1$ (note that if $|\Delta_{\mathbf{b}}| = 1$, the bargaining from time t becomes the trivial complete information bargaining) and $\iota(t) = \mathbf{b}$. Let the equilibrium offer of buyer type $\mathbf{b}_i \in \Delta_{\mathbf{b}}$ be $x_{\mathbf{b}_i}(t)$. After receiving \mathbf{b} 's offer, \mathbf{s} will update its belief and decide whether to accept the offer from \mathbf{b} . Without loss of generality, we assume that $x_{\mathbf{b}_i}(t) = -1$ if \mathbf{b}_i 's equilibrium offer will be rejected by seller \mathbf{s} at time $t+1$. There are two situations: 1) All buyer types make the same offer. In this case, a *pooling* choice rule is chosen by different buyer types. 2) Buyer

types make different offers. That is, a *separating* choice rule is used by different buyer types.

It is easy to see that there are two pooling choice rules depending on whether the seller will accept the offer at time $t + 1$ in equilibrium: 1) *accepting pooling choice rule* in which all buyer types make the same acceptable offer to seller \mathbf{s} ; 2) *rejecting pooling choice rule* in which all buyer types make the same rejectable offer (i.e., -1) to seller \mathbf{s} . While the buyer adopts the separating choice rule, some buyer types' equilibrium offers are acceptable to the seller and the number of separating choice rules is drastically reduced due to the following theorem.

Theorem 2. *There is no equilibrium assessment in pure strategies if buyer types make different acceptable offers at t .*

Proof. We can easily prove this by contradiction. Assume that there is a sequential equilibrium for a belief system in which at time t such that $\iota(t) = \mathbf{b}$, buyer \mathbf{b}_i makes an acceptable offer x to \mathbf{s} and buyer types \mathbf{b}_j makes an acceptable offer y to \mathbf{s} such that $x \neq y$. If $x > y$, buyer \mathbf{b}_i has an incentive to behave like buyer \mathbf{b}_j by offering price y . The other direction is analogous. \square

Therefore, we only need to consider the following separating choice rules: buyer types $\Delta_{\mathbf{b}}^a$ make an acceptable offer to \mathbf{s} at time t but buyer types $\Delta_{\mathbf{b}}^r = \Delta_{\mathbf{b}} \setminus \Delta_{\mathbf{b}}^a$ make an offer (i.e., -1) that will be rejected by \mathbf{s} at time t . Due to Theorem 6 (will be detailed later), we only need to consider partitions $\Delta_{\mathbf{b}}^a \cup \Delta_{\mathbf{b}}^r = \Delta_{\mathbf{b}}$ such that for any buyer type $\mathbf{b}_i \in \Delta_{\mathbf{b}}^a$ and any $\mathbf{b}_j \in \Delta_{\mathbf{b}}^r$, $RP_i > RP_j$. Thus, the number of separating choice rules is $|\Delta_{\mathbf{b}}| - 1$. For each choice rule at time t , we need to first compute the sequential equilibria for the continuation game $\Gamma(t+1)$ with corresponding *reasonable* beliefs, i.e., $\Delta_{\mathbf{b}}^a$ and $\Delta_{\mathbf{b}}^s$. Accordingly, we need to compute sequential equilibria for the continuation game with at most $2(|\Delta_{\mathbf{b}}| - 1)$ different reasonable beliefs.

Algorithm 1: Compute equilibrium strategies for a continuation game $\Gamma(t, \Delta_{\mathbf{b}})$ such that $\iota(t) = \mathbf{b}$, $|\Delta_{\mathbf{b}}| > 1$, and $t < T - 1$

Let $\mathbf{SE}(\Delta_{\mathbf{b}}, t) = \emptyset$ be the set of sequential equilibria for the continuation game with belief $\Delta_{\mathbf{b}}$ at t ;

for *each choice rule* **do**

for *each equilibrium strategy combination of the continuation game with reasonable beliefs starting from time $t + 1$* **do**

 Compute buyer types' equilibrium offers and construct offer update rules (Section 3.5);

if *equilibrium existence conditions are satisfied* **then**

 | add agents' equilibrium strategies from time t to $\mathbf{SE}(\Delta_{\mathbf{b}}, t)$;

end

end

end

return $\mathbf{SE}(\Delta_{\mathbf{b}}, t)$;

3.4.3 The Algorithm

Algorithm 1 and Algorithm 2 outline the main steps for computing agents' equilibrium strategies in a continuation game based on the sequential equilibria in its continuation game with different beliefs. To compute a buyer agent's equilibrium offer, the buyer considers different choice rules and for each choice rule, we need to consider all the sequential equilibria of the continuation game with reasonable beliefs since there may be multiple sequential equilibria for the continuation game with a specific belief. Different buyer types' equilibrium strategies are derived using a Bayesian extension of backward induction (see Section 3.5). To compute the seller's equilibrium offer at a time point, we consider all the reject update rules and for each reject update rule, we compute the sequential equilibria of the continuation game with the belief corresponding to the reject update rule. We compute the seller's equilibrium offer for each sequential equilibrium corresponding to a reject update rule and check equilibrium existence conditions (see Section 3.6).

3.4.4 Off the Equilibrium Path Optimal Strategies

Before analyzing equilibrium strategies, we provide the optimal strategies in the situations seller \mathbf{s} believes the buyer of one single type. There are two cases: 1) Seller

Algorithm 2: Compute equilibrium strategies for a continuation game $\Gamma(t, \Delta_{\mathbf{b}})$ such that $\iota(t) = \mathbf{s}$, $|\Delta_{\mathbf{b}}| > 1$, and $t < T - 1$

Let $\mathbf{SE}(\Delta_{\mathbf{b}}, t) = \emptyset$ be the set of sequential equilibria for the continuation game with belief $\Delta_{\mathbf{b}}$ at t ;

for each reject update rule **do**

for each sequential equilibrium of the continuation game with the belief corresponding to the reject update rule **do**

Compute the seller \mathbf{s} 's optimal offer and buyer types $\Delta_{\mathbf{b}}$'s acceptance decision at time $t + 1$ (Section 3.6);

if equilibrium existence conditions are satisfied **then**

| add agents' equilibrium strategies from time t to $\mathbf{SE}(\Delta_{\mathbf{b}}, t)$;

end

end

end

return $\mathbf{SE}(\Delta_{\mathbf{b}}, t)$;

\mathbf{s} has the right belief about the type of the buyer \mathbf{b} . In this case, agents' equilibrium strategies are the equilibrium strategies of the corresponding complete information bargaining discussed in Section 3.2. Let $x_{\mathbf{b}_i}^*(t)$ be any agent optimal offer at time t when \mathbf{b} is of type \mathbf{b}_i in this case. 2) Seller \mathbf{s} has the wrong belief about the type of the buyer \mathbf{b} , i.e., \mathbf{b}_i is believed to be \mathbf{b}_j .

Lemma 3. $x_{\mathbf{b}_i}^*(t) \geq x_{\mathbf{b}_j}^*(t)$ if $\text{RP}_i > \text{RP}_j$.

Proof. Case 1 ($\iota(T) = \mathbf{s}$). It follows that $x_{\mathbf{b}_i}^*(T - 1) = x_{\mathbf{b}_j}^*(T - 1) = \text{RP}_{\mathbf{s}}$. Then $x_{\mathbf{b}_i}^*(T - 2) = \text{RP}_i(1 - \delta_{\mathbf{b}}) + \delta_{\mathbf{b}}x_{\mathbf{b}_i}^*(T - 1) > x_{\mathbf{b}_j}^*(T - 2) = \text{RP}_j(1 - \delta_{\mathbf{b}}) + \delta_{\mathbf{b}}x_{\mathbf{b}_i}^*(T - 1)$. Similarly, we have $x_{\mathbf{b}_i}^*(T - 3) = \text{RP}_{\mathbf{s}}(1 - \delta_{\mathbf{s}}) + \delta_{\mathbf{s}}x_{\mathbf{b}_i}^*(T - 2)$ and $x_{\mathbf{b}_j}^*(T - 3) = \text{RP}_{\mathbf{s}}(1 - \delta_{\mathbf{s}}) + \delta_{\mathbf{s}}x_{\mathbf{b}_j}^*(T - 2)$. Thus we have $x_{\mathbf{b}_i}^*(T - 3) > x_{\mathbf{b}_j}^*(T - 3)$. Recursively, we have $x_{\mathbf{b}_i}^*(t) > x_{\mathbf{b}_j}^*(t)$ for $t < T - 3$.

Case 2 ($\iota(T) = \mathbf{b}$). It follows that $x_{\mathbf{b}_i}^*(T - 1) = \text{RP}_i > x_{\mathbf{b}_j}^*(T - 1) = \text{RP}_j$. Then at time $T - 2$, we have $x_{\mathbf{b}_i}^*(T - 2) = \text{RP}_{\mathbf{s}}(1 - \delta_{\mathbf{s}}) + \delta_{\mathbf{s}}x_{\mathbf{b}_i}^*(T - 1)$ and $x_{\mathbf{b}_j}^*(T - 2) = \text{RP}_{\mathbf{s}}(1 - \delta_{\mathbf{s}}) + \delta_{\mathbf{s}}x_{\mathbf{b}_j}^*(T - 1)$. Thus, $x_{\mathbf{b}_i}^*(T - 2) > x_{\mathbf{b}_j}^*(T - 2)$. Recursively, we have $x_{\mathbf{b}_i}^*(t) > x_{\mathbf{b}_j}^*(t)$ for $t < T - 2$. \square

We can see that \mathbf{b}_i is weaker than \mathbf{b}_j in terms of its offering price at each time point in complete information bargaining. Similarly, we can get $\text{RP}_i - x_{\mathbf{b}_i}^*(t) \geq$

$RP_j - x_{\mathbf{b}_j}^*(t)$. $RP_i - x_{\mathbf{b}_i}^*(0)$ is the gain (utility) of \mathbf{b}_i in complete information bargaining and $RP_j - x_{\mathbf{b}_j}^*(0)$ is the gain (utility) of \mathbf{b}_j in complete information bargaining.

Lemma 4. $x_{\mathbf{b}_i}^*(t) \leq (x_{\mathbf{b}_i}^*(t+1))_{\leftarrow \mathbf{b}_i}$ and $x_{\mathbf{b}_j}^*(t) \leq (x_{\mathbf{b}_j}^*(t+1))_{\leftarrow \mathbf{b}_j}$.

Proof. We can get this result by following the same procedure in the proof of Lemma 3. This result indicates that the buyer will accept sellers' lowest equilibrium price in complete information bargaining, i.e., agents will reach a final agreement at time $t - 2$ in the complete information bargaining case. \square

Agents' equilibrium strategies when seller \mathbf{s} has the wrong belief about the type of the buyer \mathbf{b} are specified in the following theorem.

Theorem 5. *If seller \mathbf{s} has the wrong belief about the type of \mathbf{b} , its optimal strategies are those in complete information bargaining. Assume that $RP_i > RP_j$. The optimal strategies $\sigma_{\mathbf{b}_i}^*(t)|\{\mathbf{b}_j\}$ of buyer \mathbf{b}_i when it is believed to be \mathbf{b}_j are:*

$$\sigma_{\mathbf{b}_i}^*(t)|\{\mathbf{b}_j\} = \begin{cases} \text{accept } y & \text{if } y \leq (x_{\mathbf{b}_j}^*(t))_{\leftarrow \mathbf{b}_i} \\ \text{offer } x_{\mathbf{b}_j}^*(t) & \text{otherwise} \end{cases}$$

The optimal strategies $\sigma_{\mathbf{b}_j}^(t)|\{\mathbf{b}_i\}$ of the buyer \mathbf{b}_j when it is believed to be \mathbf{b}_i are:*

- *If $\iota(T) = \mathbf{b}$, accept y if $y \leq \min\{(x_{\mathbf{b}_i}^*(t))_{\leftarrow \mathbf{b}_j}, RP_j\}$. Otherwise, offer $\min\{x_{\mathbf{b}_i}^*(t), RP_j\}$.*
- *If $\iota(T) = \mathbf{s}$, accept y if $y \leq \min\{(x_{\mathbf{b}_i}^*(t))_{\leftarrow \mathbf{b}_j}, (RP_j)_{\leftarrow (T-t)[\mathbf{b}_j]}\}$. Otherwise, offer $\min\{x_{\mathbf{b}_i}^*(t), (RP_{\mathbf{s}})_{\leftarrow (T-1-t)[\mathbf{b}_j]}\}$.*

Proof. Case 1 (\mathbf{b}_i is believed to be \mathbf{b}_j). If the seller offers $x_{\mathbf{b}_j}^*(t - 1)$, buyer \mathbf{b}_i 's optimal strategy is to accept it as the minimum price that the seller would accept at time $t + 1$, i.e., $x_{\mathbf{b}_j}^*(t)$, gives \mathbf{b}_i a utility lesser than $x_{\mathbf{b}_j}^*(t - 1)$ since $(x_{\mathbf{b}_j}^*(t))_{\leftarrow \mathbf{b}_i} > (x_{\mathbf{b}_j}^*(t))_{\leftarrow \mathbf{b}_j} = x_{\mathbf{b}_j}^*(t - 1)$. If the seller acts off the equilibrium path and offers a price y lower than $x_{\mathbf{b}_j}^*(t - 1)$, the optimal strategy of \mathbf{b}_i is obviously to accept y . If the seller offers a price y greater than $x_{\mathbf{b}_j}^*(t - 1)$, the optimal strategy of \mathbf{b}_i is to accept

y only if $y \leq (x_{\mathbf{b}_j}^*(t))_{\leftarrow \mathbf{b}_i}$, otherwise \mathbf{b}_i 's optimal strategy is to reject y and to offer $x_{\mathbf{b}_j}^*(t)$. Note that $x_{\mathbf{b}_i}^*(t) \leq \text{RP}_i$ and $x_{\mathbf{b}_j}^*(t) \leq \text{RP}_i$.

Case 2 (\mathbf{b}_j is believed to be \mathbf{b}_i). This case is more complicated as seller's optimal offer $x_{\mathbf{b}_i}^*(t-1)$ on its equilibrium path is not acceptable to \mathbf{b}_j as when \mathbf{b}_j offers $x_{\mathbf{b}_i}^*(t)$ at time t , $(x_{\mathbf{b}_i}^*(t))_{\leftarrow \mathbf{b}_j} < (x_{\mathbf{b}_i}^*(t))_{\leftarrow \mathbf{b}_i} = x_{\mathbf{b}_i}^*(t-1)$. In addition, \mathbf{b}_j may not offer $x_{\mathbf{b}_i}^*(t)$ if it is advantageous to wait for the agreement at time T . There are two situations: 1) $\iota(T) = \mathbf{b}$. In this case, \mathbf{s} will propose RP_i at time $T-1$, which is not acceptable to buyer \mathbf{b}_j as RP_i is higher than \mathbf{b}_j 's reserve price. Therefore, \mathbf{b}_j 's optimal offer at time t is $\min\{x_{\mathbf{b}_i}^*(t), \text{RP}_j\}$. Note that $x_{\mathbf{b}_i}^*(t)$ is always not acceptable to \mathbf{s} . 2) $\iota(T) = \mathbf{s}$. In this case, \mathbf{b}_j will propose RP_s at time $T-1$, which will be accepted by seller \mathbf{s} at time T . Therefore, \mathbf{b}_j 's optimal offer at time t is $\min\{x_{\mathbf{b}_i}^*(t), (\text{RP}_s)_{\leftarrow (T-1-t)[\mathbf{b}_j]}\}$. \square

3.5 The Buyer's Equilibrium Offer

This section focuses on computing the buyer's equilibrium offer at a continuation game $\Gamma(t, \Delta_{\mathbf{b}})$ such that $\iota(t) = \mathbf{b}$. If $t = T$, it is the buyer agent's dominant strategy to accept any offer which is not worse than its reserve price. At time $t = T-1$, different buyer types' optimal offer is RP_s since seller \mathbf{s} will accept the offer at time T . If $|\Delta_{\mathbf{b}}| = 1$, agents' equilibrium strategies are the equilibrium strategies of the corresponding complete information bargaining discussed in Section 3.2. When $|\Delta_{\mathbf{b}}| > 1$ at time $t < T-1$, buyer types have multiple choice rules and we need to consider the equilibrium strategies for each choice rule. There could be multiple equilibrium strategies for a choice rule since there could be multiple sequential equilibria for the continuation game with a reasonable belief starting from time $t+1$. In the rest of this section, we show how to compute different buyer types' equilibrium strategies given agents' equilibrium strategies of the continuation game with different beliefs and construct agents' belief systems.

Before we proceed, we introduce the concept of *equivalent offer*. In complete information bargaining, seller \mathbf{s} 's optimal offer $x^*(t)$ at time t is the value to be propagated backward at time point $t - 1$. That is, if \mathbf{b} offers $(x^*(t))_{\leftarrow \mathbf{s}}$ at time $t - 1$, \mathbf{s} will accept it at time t . With incomplete information, this property no longer holds since \mathbf{s} will accept an offer if and only if the utility of accepting the offer is not less than the *expected utility* of making its optimal offer at time t . Given the equilibrium assessment $\langle \mu, \sigma^* \rangle$, the equilibrium expected utility of seller \mathbf{s} 's offer x at time t , denoted as $EU_{\mathbf{s}}(x, t)$, is the expected utility of the seller's offering x if 1) the seller's belief at time t is $\mu(t)$ and 2) agents act according to the equilibrium strategies σ^* from time t on. The *equivalent offer* of \mathbf{s} 's offering x , denoted as $e_{\mathbf{s}}^t|\mu(t)$, is a value satisfying $U_{\mathbf{s}}(e_{\mathbf{s}}^t|\mu(t), t+1) = EU_{\mathbf{s}}(x, t)$. $e_{\mathbf{s}}^t|\mu(t)$ is the value to be propagated backward at time point $t - 1$.

Similarly, the *equivalent offer* of buyer \mathbf{b}_i 's offering x at time t , denoted as $e_{\mathbf{b}_i}^t|\mu(t)$, is a value satisfying $U_{\mathbf{b}_i}(e_{\mathbf{b}_i}^t|\mu(t), t+1) = U_{\mathbf{b}_i}(EBO(\mathbf{b}_i, x, t))$ where $EBO(\mathbf{b}_i, x, t)$ is the equilibrium bargaining outcome of \mathbf{b}_i if it offers x at time t . In addition, let $EBO(\mathbf{b}_i, \varphi)$ denote the equilibrium bargaining outcome of \mathbf{b}_i if agents follow the strategies specified by a sequential equilibrium φ . Given a bargaining outcome oc , buyer \mathbf{b}_i 's equivalent offer at time t is given by function $\rho(\mathbf{b}_i, t, oc)$ which satisfies $U_{\mathbf{b}_i}(\rho(\mathbf{b}_i, t, oc), t+1) = U_{\mathbf{b}_i}(oc)$.

3.5.1 Pooling Choice Rule

Here we consider agents' equilibrium strategies when \mathbf{b} employs a pooling choice rule at a continuation game $\Gamma(t, \Delta_{\mathbf{b}})$. Since all buyer types will behave in the same way, seller \mathbf{s} will not change its belief after observing the buyer's equilibrium offer. Thus, we need to consider all sequential equilibria $\mathbf{SE}(\Delta_{\mathbf{b}}, t+1)$ of the continuation game with belief $\Delta_{\mathbf{b}}$ at time $t+1$. If $\mathbf{SE}(\Delta_{\mathbf{b}}, t+1) = \emptyset$, there is no sequential equilibrium for this choice rule. Otherwise, for each sequential equilibrium $\varphi \in \mathbf{SE}(\Delta_{\mathbf{b}}, t+1)$,

we compute buyer types' optimal offer and check the satisfaction of equilibrium existence conditions.

First we consider the accepting pooling choice rule. Let $e_s^{t+1}|\Delta_{\mathbf{b}}$ be \mathbf{s} 's equivalent offer at time $t+1$ given the belief $\Delta_{\mathbf{b}}$ in the sequential equilibrium φ . At time $t+1$, the equilibrium strategy of \mathbf{s} is that \mathbf{s} will accept any offer y if $y \geq (e_s^{t+1}|\Delta_{\mathbf{b}})_{\leftarrow \mathbf{s}}$. Therefore, the equilibrium offer of buyer $\mathbf{b}_i \in \Delta_{\mathbf{b}}$ at time t is $(e_s^{t+1}|\Delta_{\mathbf{b}})_{\leftarrow \mathbf{s}}$.⁴ The corresponding offer update rule is the following: $\mu(t+1) = \Delta_{\mathbf{b}}$ if $\sigma_{\mathbf{b}}(t) = \text{offer}(e_s^{t+1}|\Delta_{\mathbf{b}})_{\leftarrow \mathbf{s}}$; $\mu(t+1) = \{\mathbf{b}_h(\Delta_{\mathbf{b}})\}$, otherwise.

If buyer $\mathbf{b}_i \in \Delta_{\mathbf{b}}$ deviates from offering $(e_s^{t+1}|\Delta_{\mathbf{b}})_{\leftarrow \mathbf{s}}$, it will be believed to be of type $\mathbf{b}_h(\Delta_{\mathbf{b}})$. Following Theorem 5, when a buyer \mathbf{b}_i is believed to be of type $\mathbf{b}_h(\Delta_{\mathbf{b}})$ which has a reserve price no less than RP_i , \mathbf{b}_i 's optimal offer at time t is $x_{\mathbf{b}_i}^*(t)|\{\mathbf{b}_h(\Delta_{\mathbf{b}})\}$. Thus, the condition of equilibrium existence needed to be checked is $e_{\mathbf{b}_i}^t|\Delta_{\mathbf{b}} \leq x_{\mathbf{b}_i}^*(t)|\{\mathbf{b}_h(\Delta_{\mathbf{b}})\}$ for all $\mathbf{b}_i \in \Delta_{\mathbf{b}}$. If the equilibrium existence conditions are satisfied, there is a sequential equilibrium with buyer types' offer $(e_s^{t+1}|\Delta_{\mathbf{b}})_{\leftarrow \mathbf{s}}$ and φ as the sequential equilibrium for the continuation game from time $t+1$. The sequential equilibrium will be added to $\mathbf{SE}(\Delta_{\mathbf{b}}, t)$. Buyer \mathbf{b}_i 's equilibrium bargaining outcome in this equilibrium is $EBO(\mathbf{b}_i, (e_s^{t+1}|\Delta_{\mathbf{b}})_{\leftarrow \mathbf{s}}, t) = ((e_s^{t+1}|\Delta_{\mathbf{b}})_{\leftarrow \mathbf{s}}, t+1)$ since $(e_s^{t+1}|\Delta_{\mathbf{b}})_{\leftarrow \mathbf{s}}$ is acceptable to the seller. Thus buyer \mathbf{b}_i 's equivalent offer is $e_{\mathbf{b}_i}^t|\Delta_{\mathbf{b}} = (e_s^{t+1}|\Delta_{\mathbf{b}})_{\leftarrow \mathbf{s}}$.

Then we consider the rejecting pooling choice rule. By definition, all buyer types $\Delta_{\mathbf{b}}$ will make an offer (i.e., -1) that will be rejected by the seller. Buyer \mathbf{b}_i 's equilibrium bargaining outcome is the bargaining outcome in the sequential equilibrium φ , i.e., $EBO(\mathbf{b}_i, -1, t) = EBO(\mathbf{b}_i, \varphi)$. Thus buyer \mathbf{b}_i 's equivalent offer is $e_{\mathbf{b}_i}^t|\Delta_{\mathbf{b}} = \rho(\mathbf{b}_i, t, EBO(\mathbf{b}_i, \varphi))$. We can also derive $e_{\mathbf{b}_i}^t|\Delta_{\mathbf{b}}$ in the following way. If according to the sequential equilibrium φ , seller \mathbf{s} 's equilibrium offer $x^{t+1}|\Delta_{\mathbf{b}}$ at time $t+1$ is ϖ , buyer \mathbf{b}_i 's equivalent offer at time t is then $e_{\mathbf{b}_i}^t|\Delta_{\mathbf{b}} = (e_{\mathbf{b}_i}^{t+2}|\Delta_{\mathbf{b}})_{\leftarrow 2[\mathbf{b}_i]}$ where

⁴We assume that buyer types are "cooperatively selfish" in the sense that when they are making the same acceptable offer, they will choose the lowest acceptable price.

$e_{\mathbf{b}_i}^{t+2}|\Delta_{\mathbf{b}}$ is \mathbf{b}_i 's equivalent offer at time $t + 2$. If $x^{t+1}|\Delta_{\mathbf{b}} \neq \varpi$, buyer \mathbf{b}_i 's equivalent offer at time t is 1) $e_{\mathbf{b}_i}^t|\Delta_{\mathbf{b}} = (x^{t+1}|\Delta_{\mathbf{b}})_{\leftarrow \mathbf{b}_i}$ if according to the sequential equilibrium \wp , \mathbf{s} 's equilibrium offering price $x^{t+1}|\Delta_{\mathbf{b}}$ at time $t + 1$ will be accepted by buyer type \mathbf{b}_i ; or 2) $e_{\mathbf{b}_i}^t|\Delta_{\mathbf{b}} = (e_{\mathbf{b}_i}^{t+2}|\Delta'_{\mathbf{b}})_{\leftarrow 2[\mathbf{b}_i]}$ if according to the sequential equilibrium \wp , buyer type \mathbf{b}_i will reject \mathbf{s} 's equilibrium offering price $x^{t+1}|\Delta_{\mathbf{b}}$ at time $t + 1$ where $\Delta'_{\mathbf{b}}$ is the set of buyer types that will reject \mathbf{s} 's offer at time $t + 1$. The corresponding offer update rule is the following: $\mu(t + 1) = \Delta_{\mathbf{b}}$ if $\sigma_{\mathbf{b}}(t) = offer - 1$; $\mu(t + 1) = \{\mathbf{b}_h(\Delta_{\mathbf{b}})\}$, otherwise. If buyer \mathbf{b} deviates from offering -1 at time t , it will be treated as buyer type $\mathbf{b}_h(\Delta_{\mathbf{b}})$ and the equilibrium existence condition is $e_{\mathbf{b}_i}^t|\Delta_{\mathbf{b}} \leq x_{\mathbf{b}_i}^*(t)|\{\mathbf{b}_h(\Delta_{\mathbf{b}})\}$ for all $\mathbf{b}_i \in \Delta_{\mathbf{b}}$.

3.5.2 Separating Choice Rule

Then we consider agents' equilibrium strategies at a continuation game $\Gamma(t, \Delta_{\mathbf{b}})$ when buyer \mathbf{b} employs the separating choice rule where buyer types $\Delta_{\mathbf{b}}^a$ make an acceptable offer while buyer types $\Delta_{\mathbf{b}}^r$ make a rejectable offer -1 . For this choice rule, the reasonable beliefs of its continuation game are $\Delta_{\mathbf{b}}^a$ and $\Delta_{\mathbf{b}}^r$. If one of the continuation games has no sequential equilibrium, there is no sequential equilibrium for this choice rule. We show how to compute agents' equilibrium strategies at time t given a sequential equilibrium $\wp^a \in \mathbf{SE}(\Delta_{\mathbf{b}}^a, t + 1)$ and a sequential equilibrium $\wp^r \in \mathbf{SE}(\Delta_{\mathbf{b}}^r, t + 1)$.

Let $e_{\mathbf{s}}^{t+1}|\Delta_{\mathbf{b}}^a$ be \mathbf{s} 's equivalent offer at time $t + 1$ in the equilibrium \wp^a . Let $e_{\mathbf{s}}^{t+1}|\Delta_{\mathbf{b}}^r$ ($x^{t+1}|\Delta_{\mathbf{b}}^r$, respectively) be \mathbf{s} 's equivalent offer (equilibrium offer, respectively) at time $t + 1$ in the equilibrium \wp^r . Similar to the pooling acceptance choice rule, the optimal offer of buyer types $\Delta_{\mathbf{b}}^a$ at time t is $(e_{\mathbf{s}}^{t+1}|\Delta_{\mathbf{b}}^a)_{\leftarrow \mathbf{s}}$. Accordingly, buyer $\mathbf{b}_i \in \Delta_{\mathbf{b}}^a$'s equivalent offer is $e_{\mathbf{b}_i}^t|\Delta_{\mathbf{b}} = (e_{\mathbf{s}}^{t+1}|\Delta_{\mathbf{b}}^a)_{\leftarrow \mathbf{s}}$ since its equilibrium bargaining outcome is $((e_{\mathbf{s}}^{t+1}|\Delta_{\mathbf{b}}^a)_{\leftarrow \mathbf{s}}, t + 1)$. By convention, the equilibrium offer of buyer type $\mathbf{b}_j \in \Delta_{\mathbf{b}}^r$ at time t is -1 . Buyer \mathbf{b}_j 's equilibrium bargaining outcome is the bargaining outcome

$EBO(\mathbf{b}_j, \wp^r)$ in the sequential equilibrium \wp^r . Thus buyer $\mathbf{b}_j \in \Delta_{\mathbf{b}}^r$'s equivalent offer is $e_{\mathbf{b}_j}^t | \Delta_{\mathbf{b}} = \rho(\mathbf{b}_j, t, EBO(\mathbf{b}_j, \wp^r))$. We can also compute $e_{\mathbf{b}_j}^t | \Delta_{\mathbf{b}}$ by considering the following two situations:

- If $x^{t+1} | \Delta_{\mathbf{b}}^r \neq \varpi$, buyer $\mathbf{b}_j \in \Delta_{\mathbf{b}}^r$'s equivalent offer at time t is 1) $e_{\mathbf{b}_j}^t | \Delta_{\mathbf{b}} = (x^{t+1} | \Delta_{\mathbf{b}}^r)_{\leftarrow \mathbf{b}_j}$ if according to the sequential equilibrium \wp^r , \mathbf{s} 's equilibrium offering price $x^{t+1} | \Delta_{\mathbf{b}}^r$ at time $t + 1$ will be accepted by buyer type \mathbf{b}_j ; or 2) $e_{\mathbf{b}_j}^t | \Delta_{\mathbf{b}} = (e_{\mathbf{b}_j}^{t+2} | \Delta_{\mathbf{b}}^{r'})_{\leftarrow 2[\mathbf{b}_j]}$ if according to the sequential equilibrium \wp^r , buyer type \mathbf{b}_j will reject \mathbf{s} 's equilibrium offering price $x^{t+1} | \Delta_{\mathbf{b}}^r$ at time $t + 1$ where $\Delta_{\mathbf{b}}^{r'} \subseteq \Delta_{\mathbf{b}}^r$ is the set of buyer types that will reject \mathbf{s} 's offer at time $t + 1$.
- If $x^{t+1} | \Delta_{\mathbf{b}}^r = \varpi$, buyer $\mathbf{b}_j \in \Delta_{\mathbf{b}}^r$'s equivalent offer at time t is then $e_{\mathbf{b}_j}^t | \Delta_{\mathbf{b}} = (e_{\mathbf{b}_j}^{t+2} | \Delta_{\mathbf{b}}^r)_{\leftarrow 2[\mathbf{b}_j]}$ where $e_{\mathbf{b}_j}^{t+2} | \Delta_{\mathbf{b}}^r$ is \mathbf{b}_j 's equivalent offer at time $t + 2$.

Seller \mathbf{s} will update its belief to $\Delta_{\mathbf{b}}^a$ when it receives an offer $(e_{\mathbf{s}}^{t+1} | \Delta_{\mathbf{b}}^a)_{\leftarrow \mathbf{s}}$. If it receives an offer -1 , it will update its belief to $\Delta_{\mathbf{b}}^r$. Otherwise, it will update its belief to $\mathbf{b}_h(\Delta_{\mathbf{b}})$. The existence of such an equilibrium depends on the following conditions:

- Any buyer type $\mathbf{b}_i \in \Delta_{\mathbf{b}}^a$ has no incentive to behave as any buyer type $\mathbf{b}_j \in \Delta_{\mathbf{b}}^r$. If \mathbf{b}_i pretends to be \mathbf{b}_j , it will offer -1 at time t and its equilibrium bargaining outcome will be $EBO(\mathbf{b}_j, -1, t) = EBO(\mathbf{b}_j, \wp^r)$. Therefore, this condition requires that $U_{\mathbf{b}_i}(EBO(\mathbf{b}_i, (e_{\mathbf{s}}^{t+1} | \Delta_{\mathbf{b}}^a)_{\leftarrow \mathbf{s}}, t)) \geq U_{\mathbf{b}_i}(EBO(\mathbf{b}_j, \wp^r))$ or equivalently, $(e_{\mathbf{s}}^{t+1} | \Delta_{\mathbf{b}}^a)_{\leftarrow \mathbf{s}} \leq \rho(\mathbf{b}_i, t, EBO(\mathbf{b}_j, \wp^r))$.
- Any buyer type $\mathbf{b}_j \in \Delta_{\mathbf{b}}^r$ must have no incentive to behave as $\mathbf{b}_i \in \Delta_{\mathbf{b}}^a$. If \mathbf{b}_j behaves as \mathbf{b}_i , it will offer $(e_{\mathbf{s}}^{t+1} | \Delta_{\mathbf{b}}^a)_{\leftarrow \mathbf{s}}$ at time t and the offer will be accepted. \mathbf{b}_j will not choose to behave as \mathbf{b}_i if $U_{\mathbf{b}_j}(EBO(\mathbf{b}_j, \wp^r)) \geq U_{\mathbf{b}_j}(EBO(\mathbf{b}_j, (e_{\mathbf{s}}^{t+1} | \Delta_{\mathbf{b}}^a)_{\leftarrow \mathbf{s}}, t))$ or equivalently, $\rho(\mathbf{b}_j, t, EBO(\mathbf{b}_j, \wp^r)) \leq (e_{\mathbf{s}}^{t+1} | \Delta_{\mathbf{b}}^a)_{\leftarrow \mathbf{s}}$.
- No buyer type has an incentive to offer a price different from the above two equilibrium offers. If a buyer type $\mathbf{b}_i \in \Delta_{\mathbf{b}}$ offers a price different from $(e_{\mathbf{s}}^{t+1} | \Delta_{\mathbf{b}}^a)_{\leftarrow \mathbf{s}}$

and -1 , it will be treated as buyer type $\mathbf{b}_h(\Delta_{\mathbf{b}})$ and its optimal offer at time t is then $x_{\mathbf{b}_i}^*(t) | \{\mathbf{b}_h(\Delta_{\mathbf{b}})\}$. Buyer type \mathbf{b}_i will not choose to act off the equilibrium path if $e_{\mathbf{b}_i}^t | \Delta_{\mathbf{b}} \leq x_{\mathbf{b}_i}^*(t) | \{\mathbf{b}_h(\Delta_{\mathbf{b}})\}$.

If all the three conditions are satisfied, buyer types' optimal offers, the belief update rule, and the sequential equilibria \wp^a and \wp^r for the continuation game starting from time $t+1$ consists of a sequential equilibrium for the continuation game $\Gamma(t, \Delta_{\mathbf{b}})$. The following theorem suggests that we only need to consider at most $|\Delta_{\mathbf{b}}|$ different choice rules.

Theorem 6. *Assume that \mathbf{b} behaves in different ways at a continuation game with belief set $\Delta_{\mathbf{b}}$ where $\Delta_{\mathbf{b}} = \Delta_{\mathbf{b}}^a \cup \Delta_{\mathbf{b}}^r$ at time t . If there is a buyer type $\mathbf{b}_i \in \Delta_{\mathbf{b}}^a$ and a buyer $\mathbf{b}_j \in \Delta_{\mathbf{b}}^r$ such that $\text{RP}_i < \text{RP}_j$, there is no sequential equilibrium for this choice rule.*

Proof. This result can be proved by contradiction. If there is a sequential equilibrium, the following two conditions are satisfied: 1) Buyer type \mathbf{b}_i has no incentive to behave as \mathbf{b}_j , i.e., $(e_{\mathbf{s}}^{t+1} | \Delta_{\mathbf{b}}^a)_{\leftarrow \mathbf{s}} \leq \rho(\mathbf{b}_i, t, \text{EBO}(\mathbf{b}_j, \wp^r))$; and 2) Buyer type \mathbf{b}_j has no incentive to behave as \mathbf{b}_i , i.e., $\rho(\mathbf{b}_j, t, \text{EBO}(\mathbf{b}_j, \wp^r)) \leq (e_{\mathbf{s}}^{t+1} | \Delta_{\mathbf{b}}^a)_{\leftarrow \mathbf{s}}$. Therefore, equilibrium existence requires that $\rho(\mathbf{b}_j, t, \text{EBO}(\mathbf{b}_j, \wp^r)) \leq \rho(\mathbf{b}_i, t, \text{EBO}(\mathbf{b}_j, \wp^r))$.

Assume that $\text{EBO}(\mathbf{b}_j, \wp^r) = (x, t')$ where $T \geq t' > t$. From the definition of equivalent offers, we have $(\text{RP}_j - \rho(\mathbf{b}_j, t, \text{EBO}(\mathbf{b}_j, \wp^r))) \cdot \delta_{\mathbf{b}}^{t+1} = (\text{RP}_j - x) \cdot \delta_{\mathbf{b}}^{t'}$, which can be rewritten as $\rho(\mathbf{b}_j, t, \text{EBO}(\mathbf{b}_j, \wp^r)) = \text{RP}_j - (\text{RP}_j - x) \cdot \delta_{\mathbf{b}}^{t'-t-1}$. Similarly, we have $\rho(\mathbf{b}_i, t, \text{EBO}(\mathbf{b}_j, \wp^r)) = \text{RP}_i - (\text{RP}_i - x) \cdot \delta_{\mathbf{b}}^{t'-t-1}$. Since $\text{RP}_i < \text{RP}_j$, it follows that $\rho(\mathbf{b}_i, t, \text{EBO}(\mathbf{b}_j, \wp^r)) < \rho(\mathbf{b}_j, t, \text{EBO}(\mathbf{b}_j, \wp^r))$ which contradicts with equilibrium existence conditions. \square

This result drastically reduces the number of separating choice rules we need to consider. Consider an belief set $\Delta_{\mathbf{b}}$ at time $t < T-1$ such that $|\Delta_{\mathbf{b}}| > 1$ and $\iota(t) = \mathbf{b}$. The total number of partitions satisfying the condition $\Delta_{\mathbf{b}}^a \cup \Delta_{\mathbf{b}}^r = \Delta_{\mathbf{b}}$ is $2^{|\Delta_{\mathbf{b}}|} - 2$.

Given Theorem 6, we just need to consider partitions $\Delta_{\mathbf{b}}^a \cup \Delta_{\mathbf{b}}^r = \Delta_{\mathbf{b}}$ such that for any buyer type $\mathbf{b}_i \in \Delta_{\mathbf{b}}^a$ and any $\mathbf{b}_j \in \Delta_{\mathbf{b}}^r$, $\text{RP}_i > \text{RP}_j$. Then for each belief set $\Delta_{\mathbf{b}}$, the total number of separating choice rules is $|\Delta_{\mathbf{b}}| - 1$.

3.6 The Seller's Equilibrium Offer

This section discusses how to compute the seller's equilibrium offer at a continuation game $\Gamma(t, \Delta_{\mathbf{b}})$ such that $\iota(t) = \mathbf{s}$. If $t = T$, it is the seller's dominant strategy to accept any offer which is not worse than its reserve price. At time $t = T - 1$, seller \mathbf{s} has multiple choices, each for one buyer type in $\Delta_{\mathbf{b}}$. The optimal offer of seller \mathbf{s} for buyer type $\mathbf{b}_i \in \Delta_{\mathbf{b}}$ is RP_i , which gives seller \mathbf{s} an expected utility $EU_{\mathbf{s}}(\text{RP}_i, T - 1) = \sum_{\mathbf{b}_j \in \Delta_{\mathbf{b}}, \text{RP}_j \geq \text{RP}_i} \omega_{\mathbf{b}_j}(\Delta_{\mathbf{b}}) U_{\mathbf{s}}(\text{RP}_i, T)$ since RP_i is only acceptable to a buyer type with a reserve price no less than RP_i . The optimal offer of \mathbf{s} at time $T - 1$ is $y = \arg \max_{y \in \{\text{RP}_i | \mathbf{b}_i \in \Delta_{\mathbf{b}}\}} EU_{\mathbf{s}}(y, T - 1)$ and its equivalent offer is $e_{\mathbf{s}}^{T-1} | \Delta_{\mathbf{b}}$ such that $U_{\mathbf{s}}(e_{\mathbf{s}}^{T-1} | \Delta_{\mathbf{b}}, T) = EU_{\mathbf{s}}(y, T - 1)$. If $|\Delta_{\mathbf{b}}| = 1$, agents' equilibrium strategies are the equilibrium strategies of the corresponding complete information bargaining discussed in Section 3.2.

Our idea of computing the seller's equilibrium offer given a belief $\Delta_{\mathbf{b}}$ ($|\Delta_{\mathbf{b}}| > 1$) at time $t < T - 1$ is the following. We consider all possible reject update rules and for each reject update rule, we compute all the sequential equilibria for the continuation game with beliefs corresponding to the reject update rule. Then for each sequential equilibrium for a reject update rule, we compute the seller's optimal offer and check the equilibrium existence conditions. Theorem 1 suggests that we only need to consider reject update rules in which buyer types with higher reserve prices accept the seller's equilibrium offer while buyer types with lower reserve prices reject the seller's equilibrium offer. Thus we only need to consider a restricted set of beliefs for the continuation game.

We consider a reject update rule in which buyer types Δ'_b will reject the seller's offer and buyer types $\Delta_b - \Delta'_b$ will accept the seller's offer such that $RP_i > RP_j$ for any $\mathbf{b}_i \in \Delta_b - \Delta'_b$ and $\mathbf{b}_j \in \Delta'_b$. We first compute all the sequential equilibria $\mathbf{SE}(\Delta'_b, t+1)$ for the continuation game with belief Δ'_b starting from time $t+1$. If there is no sequential equilibrium for the continuation game $\Gamma(t+1, \Delta'_b)$, there is no sequential equilibrium for this reject update rule. Otherwise, for each sequential equilibrium $\varphi \in \mathbf{SE}(\Delta'_b, t+1)$, we check whether there exists a price x such that the price, the reject update rule, and the sequential equilibrium φ constitute a sequential equilibrium for the continuation game $\Gamma(t, \Delta_b)$.

A price x , a reject update rule, and a sequential equilibrium φ constitute a sequential equilibrium if and only if the following three conditions are satisfied (assume that $\mathbf{b}_i \in \Delta_b - \Delta'_b$ and $\mathbf{b}_j \in \Delta'_b$ where buyer types Δ'_b will reject the seller's offer x):

1. \mathbf{b}_i is willing to accept the offer x and does not want to behave as \mathbf{b}_j . That is, for any $\mathbf{b}_i \in \Delta_b - \Delta'_b$ and $\mathbf{b}_j \in \Delta'_b$, $U_{\mathbf{b}_i}(x, t+1) \geq U_{\mathbf{b}_i}(EBO(\mathbf{b}_j, \varphi))$ where $EBO(\mathbf{b}_j, \varphi)$ is the \mathbf{b}_i 's equilibrium bargaining outcome when it behaves as \mathbf{b}_j . This condition can be reformulated as $x \leq \min_{\mathbf{b}_i \in \Delta_b - \Delta'_b, \mathbf{b}_j \in \Delta'_b} \rho(\mathbf{b}_i, t, EBO(\mathbf{b}_j, \varphi))$, which provides an upper bound for seller's offering price x . Intuitively, if the offering price x is too high (e.g., higher than RP_i), \mathbf{b}_i can not accept the offering price.
2. \mathbf{b}_j will reject the offer x . That is, each buyer type $\mathbf{b}_j \in \Delta'_b$ has no incentive to behave as \mathbf{b}_i , i.e., $U_{\mathbf{b}_j}(x, t+1) < U_{\mathbf{b}_j}(EBO(\mathbf{b}_j, \varphi))$. This condition can be rewritten as $x > \max_{\mathbf{b}_j \in \Delta'_b} \rho(\mathbf{b}_j, t, EBO(\mathbf{b}_j, \varphi))$, which provides a lower bound for the offering price x . Intuitively, if the offering price x is very low (e.g., close to 0), \mathbf{b}_j will choose to accept the favorite offer.
3. Seller \mathbf{s} has no incentive to choose a price other than x given the reject update rule and the sequential equilibrium φ of the continuation game $\Gamma(t+1, \Delta'_b)$;

The third condition requires that the price x is seller's optimal offer given the reject update rule and the sequential equilibrium φ for the continuation game. Any buyer type can either accept the seller's offer x or reject it and receive a bargaining outcome in the sequential equilibrium φ for the continuation game. Formally, buyer type $\mathbf{b}_j \in \Delta'_b$ will accept a price x if and only if $x \leq \rho(\mathbf{b}_j, t, EBO(\mathbf{b}_j, \varphi))$. Buyer type $\mathbf{b}_i \in \Delta_b - \Delta'_b$ will accept a price x if and only if $x \leq \min_{\mathbf{b}_j \in \Delta'_b} \rho(\mathbf{b}_i, t, EBO(\mathbf{b}_j, \varphi))$. We can define the acceptance price $\phi(\mathbf{b}_i, \Delta'_b, \varphi)$ of each buyer type $\mathbf{b}_i \in \Delta_b$ given the sequential equilibrium φ as follows:

$$\phi(\mathbf{b}_i, \Delta'_b, \varphi) = \begin{cases} \rho(\mathbf{b}_i, t, EBO(\mathbf{b}_i, \varphi)) & \text{if } \mathbf{b}_i \in \Delta'_b \\ \min_{\mathbf{b}_j \in \Delta'_b} \rho(\mathbf{b}_i, t, EBO(\mathbf{b}_j, \varphi)) & \text{otherwise} \end{cases}$$

Seller \mathbf{s} 's expected utility of making an offer x given the sequential equilibrium φ is defined as

$$EU_s(x, t) = \sum_{\mathbf{b}_i \in \Delta_b} \omega_{\mathbf{b}_i}(\Delta_b) EU_s(x, t, \mathbf{b}_i)$$

where $EU_s(x, t, \mathbf{b}_i)$ is seller \mathbf{s} 's utility if the buyer is of type \mathbf{b}_i , which is defined as

$$\begin{cases} U_s(x, t+1) & \text{if } x \leq \phi(\mathbf{b}_i, \Delta'_b, \varphi) \\ U_s(EBO(\mathbf{b}_i, \varphi)) & \text{if } x > \phi(\mathbf{b}_i, \Delta'_b, \varphi) \text{ and } \mathbf{b}_i \in \Delta'_b \\ U_s(\min_{\mathbf{b}_j \in \Delta'_b} \rho(\mathbf{b}_i, t, EBO(\mathbf{b}_j, \varphi)), t+1) & \text{otherwise} \end{cases}$$

It is easy to see that the optimal offer the seller should be either one buyer type's acceptance price or a price that will be rejected by all buyer types (i.e., ϖ). If the seller's optimal offer x satisfies the first two equilibrium existence conditions, there is a sequential equilibrium in which the seller offers price x and buyer types Δ'_b will reject the offer with the sequential equilibrium φ . If such a x value does not exist,

there is no sequential equilibrium given this reject update rule and the continuation game equilibrium \wp .

In addition to the above reject update rules under which at least one buyer type will choose to reject the offer, the seller can also make an offer such that it is all buyer types' equilibrium strategies to accept the offer. It is easy to see that the highest offer that will be accepted by all buyer types in equilibrium is $x = \min_{\mathbf{b}_i \in \Delta_{\mathbf{b}}} (x_{\mathbf{b}_i}^*(t + 1) | \mathbf{b}_h(\Delta_{\mathbf{b}}))_{\leftarrow \mathbf{b}_i}$ since if a seller offers a price larger than x , at least one buyer type has an incentive to deviate from accepting the offer. If the buyer rejects x , the seller will update its belief to $\mathbf{b}_h(\Delta_{\mathbf{b}})$. The acceptance price of buyer type \mathbf{b}_i for this reject update rule is thus $(x_{\mathbf{b}_i}^*(t + 1) | \mathbf{b}_h(\Delta_{\mathbf{b}}))_{\leftarrow \mathbf{b}_i}$. If the optimal offer of the seller in this case is not acceptable to all the buyer types (i.e., the optimal offer is not $\min_{\mathbf{b}_i \in \Delta_{\mathbf{b}}} (x_{\mathbf{b}_i}^*(t + 1) | \mathbf{b}_h(\Delta_{\mathbf{b}}))_{\leftarrow \mathbf{b}_i}$), there is no sequential equilibrium for this null reject update rule. Otherwise, there is a sequential equilibrium in which the seller will make an offer which will be accepted by all buyer types.

3.7 Equilibrium Existence

The algorithm for producing equilibrium strategies is a backward induction process, which starts from the continuation game with the initial belief at time $t = 0$. Here we show an example of equilibrium calculation for the bargaining game with the following parameters: $T = 3$, $\iota(0) = \mathbf{s}$, $\text{RP}_{\mathbf{s}} = 10$, $\text{RP}_1 = 100$, $\text{RP}_2 = 60$, $\text{RP}_3 = 50$, $\omega_{\mathbf{b}_1}^0 = 0.25$, $\omega_{\mathbf{b}_2}^0 = 0.5$, $\omega_{\mathbf{b}_3}^0 = 0.25$, $\delta_{\mathbf{s}} = 0.8$, and $\delta_{\mathbf{b}} = 0.6$. Before we start the backward induction process, we compute agents' equilibrium offers in complete information setting using the approach in Section 3.2 since we may use these equilibrium offers to construct agents' equilibrium strategies for the bargaining game. If buyer is of type \mathbf{b}_1 , we have $x_{\mathbf{b}_1}^*(2) = 100$, $x_{\mathbf{b}_1}^*(1) = 82$, and $x_{\mathbf{b}_1}^*(0) = 89.2$. If buyer is of type \mathbf{b}_2 , we have $x_{\mathbf{b}_2}^*(2) = 60$, $x_{\mathbf{b}_2}^*(1) = 50.0$, and $x_{\mathbf{b}_2}^*(0) = 54.0$. If buyer is of type \mathbf{b}_3 , we have $x_{\mathbf{b}_3}^*(2) = 50$, $x_{\mathbf{b}_3}^*(1) = 42.0$, and $x_{\mathbf{b}_3}^*(0) = 45.2$.

Now we try our algorithm to compute all the sequential equilibria. Since $\iota(0) = \mathbf{s}$, seller \mathbf{s} will consider different reject update rules at time $t = 0$. Therefore, we need to first compute all the sequential equilibria for the continuation game with reasonable beliefs $\{\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3\}$, $\{\mathbf{b}_2, \mathbf{b}_3\}$, and $\{\mathbf{b}_3\}$ at time $t = 1$. We show how to compute sequential equilibria for the continuation game with belief $\{\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3\}$ at time $t = 1$. At time $t = 1$, \mathbf{b} can apply different choice rules and we need to first compute all the sequential equilibria for the continuation game with beliefs $\{\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3\}$, $\{\mathbf{b}_1\}$, $\{\mathbf{b}_2, \mathbf{b}_3\}$, $\{\mathbf{b}_1, \mathbf{b}_2\}$, and $\{\mathbf{b}_3\}$ at time $t = 2$. For the continuation game with belief $\{\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3\}$ at time $t = 2$, \mathbf{s} can offer $\text{RP}_3 = 50$, $\text{RP}_2 = 60$, or $\text{RP}_1 = 100$: If it offers 50, its expected utility is $(50 - 10)0.8^3 = 20.48$; If it offers 60, its expected utility is $0.75(60 - 10)0.8^3 = 19.20$; If it offers 100, its expected utility is $0.25(100 - 10)0.8^3 = 11.52$. Thus, the optimal offer of \mathbf{s} at time $t = 3$ is 50 and the equivalent price is $e_s^2|\{\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3\} = 50.0$. When the belief is $\{\mathbf{b}_2, \mathbf{b}_3\}$, the optimal offer of \mathbf{s} at time $t = 3$ is 50 and the equivalent price is $e_s^2|\{\mathbf{b}_2, \mathbf{b}_3\} = 50.0$. When the belief is $\{\mathbf{b}_1, \mathbf{b}_2\}$, the optimal offer of \mathbf{s} at time $t = 3$ is 60 and the equivalent price is $e_s^2|\{\mathbf{b}_1, \mathbf{b}_2\} = 60.0$. For the pooling accepting choice rule at time $t = 1$, the optimal offer of all buyer types is $(e_s^2|\{\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3\})_{\leftarrow \mathbf{s}} = (50)_{\leftarrow \mathbf{s}} = 42$ and no buyer has an incentive to deviate from it: if any buyer type chooses a different offer, it will be treated as buyer type \mathbf{b}_1 and its utility in the later negotiation is 0. For the pooling rejecting choice rule at time $t = 1$, all buyer types will offer -1 and no buyer has an incentive to deviate from it. For example, buyer \mathbf{b}_2 's equivalent offer of offering -1 is $e_{\mathbf{b}_2}^1|\{\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3\} = (e_s^2|\{\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3\})_{\leftarrow \mathbf{b}_2} = (50)_{\leftarrow \mathbf{b}_2} = 54$. If buyer \mathbf{b}_2 deviates from offering -1 , it will be believed to be \mathbf{b}_1 and its optimal offer at time $t = 1$ is then $x_{\mathbf{b}_2}^*(1)|\{\mathbf{b}_1\} = 60$, which is higher than $e_{\mathbf{b}_2}^1|\{\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3\}$. For the separating choice rule in which \mathbf{b}_1 makes an acceptable offer while the other two buyer types offer -1 , there is only one sequential equilibrium for the continuation game with beliefs $\{\mathbf{b}_1\}$ and $\{\mathbf{b}_2, \mathbf{b}_3\}$. Thus, the optimal offer of \mathbf{b}_1 is $x_{\mathbf{b}_1}^1|\{\mathbf{b}_1\} = x_{\mathbf{b}_1}^*(1) = 82$. However,

\mathbf{b}_1 has an incentive to behave as \mathbf{b}_2 or \mathbf{b}_3 since in the sequential equilibrium for the continuation game with beliefs $\{\mathbf{b}_2, \mathbf{b}_3\}$, \mathbf{s} will make an offer 50 at time $t = 2$ which can bring \mathbf{b}_1 a higher utility because $(50)_{\leftarrow \mathbf{b}_1} = 70 < 82$. Therefore, there is no sequential equilibrium for this choice rule. However we can show that there is a sequential equilibrium for the separating choice rule in which \mathbf{b}_1 and \mathbf{b}_2 make an acceptable offer and their optimal offer is $(e_s^2|\{\mathbf{b}_1, \mathbf{b}_2\})_{\leftarrow \mathbf{s}} = (60)_{\leftarrow \mathbf{s}} = 50$. There are totally 3 sequential equilibria for the continuation game with belief $\{\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3\}$ at time $t = 1$.

In the same way, we can compute the two sequential equilibria for the continuation game with belief $\{\mathbf{b}_2, \mathbf{b}_3\}$ at time $t = 1$. One equilibrium is for the pooling accepting choice rule in which both buyer types offer 42. In this case, both buyer types will accept the offer in equilibrium. The other is for the separating choice rule in which buyer \mathbf{b}_2 offers 50 but buyer \mathbf{b}_3 offers -1 . In this case, \mathbf{b}_2 's offer will be accepted at time $t = 2$. \mathbf{b}_3 's offer will be rejected and the seller will offer 50 after updating its belief to $\{\mathbf{b}_3\}$.

Now we consider the seller's equilibrium offers with the initial belief at time $t = 0$. Seller \mathbf{s} can apply the following different reject update rules:

1. If the buyer rejects the seller \mathbf{s} 's offer, seller \mathbf{s} will not change its belief, i.e., seller \mathbf{s} is offering ϖ by convention. Under this reject update rule, any buyer type's acceptance price depends on the negotiation outcome in the continuation game $\Gamma(1, \{\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3\})$, which has three sequential equilibria. In the first sequential equilibrium where all buyer types adopt the pooling accepting choice rule at time $t = 1$, all buyer types will offer 42 and the seller will accept it at time $t = 2$. If seller \mathbf{s} offers ϖ , its utility is $(42 - 10)0.8^2 = 20.48$. Buyer types' acceptance prices are $(42)_{\leftarrow \mathbf{b}_1} = 65.2$ for \mathbf{b}_1 , $(42)_{\leftarrow \mathbf{b}_2} = 49.2$ for \mathbf{b}_2 , $(42)_{\leftarrow \mathbf{b}_3} = 45.2$ for \mathbf{b}_3 . If seller \mathbf{s} offers 65.2, buyer type \mathbf{b}_1 will accept it but buyer types \mathbf{b}_2 and \mathbf{b}_3 will reject it and follow the first sequential equilibrium. Thus, seller \mathbf{s} 's expected

utility while offering 65.2 is $0.25 * (65.2 - 10) * 0.8 + 0.75 * (42 - 10)0.8^2 = 26.40$. In the same way, we can find that the seller achieves an expected utility of 28.64 while offering 49.2 and achieves an expected utility of 28.16 while offering 45.2. Thus, the seller's optimal offer is 49.2. However, buyer types \mathbf{b}_1 and \mathbf{b}_2 will accept the offer, which is in conflict with the reject update rule in which all buyer types will reject the offer. Therefore, there is no sequential equilibrium for the first sequential equilibrium for the continuation game $\Gamma(1, \{\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3\})$. In the same way, we can see that there is no sequential equilibrium for the other two sequential equilibria for the continuation game $\Gamma(1, \{\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3\})$.

2. If the buyer rejects the seller \mathbf{s} 's offer, seller \mathbf{s} updates its belief to $\{\mathbf{b}_3\}$, i.e., seller \mathbf{s} is making an offer acceptable to $\mathbf{b}_1, \mathbf{b}_2$. There is only one sequential equilibrium for the continuation game $\Gamma(1, \{\mathbf{b}_3\})$ in which buyer type \mathbf{b}_3 offers $x_{\mathbf{b}_3}^*(1) = 42.0$. It is easy to see that seller \mathbf{s} gets a utility of $(42 - 10)0.8^2 = 20.48$ if it offers ϖ at time 0. Buyer types' acceptance prices are $(42)_{\leftarrow \mathbf{b}_1} = 65.2$ for \mathbf{b}_1 , $(42)_{\leftarrow \mathbf{b}_2} = 49.2$ for \mathbf{b}_2 , $(42)_{\leftarrow \mathbf{b}_3} = 45.2$ for \mathbf{b}_3 . Seller \mathbf{s} 's optimal offer for this reject update rule is 49.2. We can easily see that no buyer type has an incentive to deviate: 1) By rejecting the offer 49.2, buyer type \mathbf{b}_1 can gain a utility of $(100 - 42)0.6^2 = 20.88$, which is lower than the utility $(100 - 49.2)0.6 = 30.48$ when it accepts the offer. 2) Buyer type \mathbf{b}_2 gains a utility of $(60 - 49.2)0.6 = 6.48$ by accepting the offer 49.2, which is not lower than its utility $(60 - 42)0.6^2 = 6.48$ when it rejects the offer. 3) Buyer type \mathbf{b}_3 gains a utility of $(50 - 49.2)0.6 = 0.48$ by accepting the offer 49.2, which is lower than its utility $(50 - 42)0.6^2 = 2.88$ when it rejects the offer.
3. If the buyer rejects the seller \mathbf{s} 's offer, seller \mathbf{s} updates its belief to $\{\mathbf{b}_2, \mathbf{b}_3\}$, i.e., seller \mathbf{s} is making an offer only acceptable to \mathbf{b}_1 . There are two sequential equilibria for the continuation game $\Gamma(1, \{\mathbf{b}_2, \mathbf{b}_3\})$. For the sequential equilib-

rium in which both buyer types offer an acceptable price 42, the seller's optimal offer is 49.2 such that \mathbf{b}_1 will accept it and \mathbf{b}_2 and \mathbf{b}_3 will reject it. However, we could not find an offer which satisfies all equilibrium existence conditions given the other sequential equilibrium for the continuation game $\Gamma(1, \{\mathbf{b}_2, \mathbf{b}_3\})$.

4. Finally, the seller can make an offer that will be accepted by all buyer types, i.e., seller \mathbf{s} updates its belief to $\{\mathbf{b}_1\}$ if buyer rejects the seller \mathbf{s} 's offer. The acceptance prices for buyer types \mathbf{b}_1 , \mathbf{b}_2 , and \mathbf{b}_3 are 89.2, 60.0, and 50.0, respectively. Seller \mathbf{s} 's optimal offer for this reject update rule is 89.2. Buyer types \mathbf{b}_2 and \mathbf{b}_3 will reject the seller's optimal offer, which is in conflicting with the reject update rule. Therefore, there is no sequential equilibrium for this reject update rule.

Therefore, there are two sequential equilibria for the bargaining game. In the first equilibrium, seller \mathbf{s} will offer 49.2 at time $t = 0$. If the buyer is of type \mathbf{b}_1 or \mathbf{b}_2 , it will accept the offer and they make an agreement at time $t = 1$. If the buyer is of type \mathbf{b}_3 , it will reject the offer and make a counter offer 42.0 at time $t = 1$. When seller \mathbf{s} receives offer 42.0, it will update its belief to $\{\mathbf{b}_3\}$ and it will accept the offer at time $t = 2$. In the second equilibrium, seller \mathbf{s} will also offer 49.2 at time $t = 0$. If the buyer is of type \mathbf{b}_1 , it will accept the offer and they make an agreement at time $t = 1$. If the buyer is of type \mathbf{b}_2 or \mathbf{b}_3 , it will reject the offer and make a counter offer 42.0 at time $t = 1$. When seller \mathbf{s} receives offer 42.0, it will update its belief to $\{\mathbf{b}_2, \mathbf{b}_3\}$ and it will accept the offer at time $t = 2$.

The following theorem states that the proposed approach is sound and is complete.

Theorem 7. *Our algorithm can generate all pure strategy sequential equilibria.*

Proof. Our algorithm is complete since at any decision making point, we consider 1) all sequential equilibria of the continuation game with different beliefs, 2) all choice

rules when it's the buyer's turn to make an offer, and 3) all possible reject update rules if it's the seller's turn to make an offer.

If all equilibrium existence conditions are satisfied, agents' strategies and belief systems generated by our algorithm constitutes of a sequential equilibrium. The sequential rationality is easily seen from the backward construction: agents' strategies at time t is optimal in the continuation game starting from time t . Consistency can be proved by the assessment sequence $a_n = (\mu_n, \sigma_n)$ where σ_n is the fully mixed strategy profile such that for the seller and buyer type $\mathbf{b}_h(\Delta_{\mathbf{b}}^0)$ there is probability $1 - 1/n$ of performing the action prescribed by the equilibrium strategy profile and the remaining probability $1/n$ is uniformly distributed among the other allowed actions, while for any other buyer type $\mathbf{b}_i \in \Delta_{\mathbf{b}}^0 - \mathbf{b}_h(\Delta_{\mathbf{b}}^0)$, there is probability $1 - 1/n^T$ of performing the action prescribed by the equilibrium strategy profile and the remaining probability $1/n^T$ is uniformly distributed among the other allowed actions, and μ_n is the system of beliefs obtained applying Bayes rule starting from the same *priori* probability distribution $P_{\mathbf{b}}^0$. As $n \rightarrow \infty$, the above mixed strategy profile converges to the equilibrium strategy profile. In addition, the beliefs generated by the mixed strategy profile converges to the *priori* probability distribution. Thus, the assessment is consistent. \square

Since we focus on pure strategy equilibrium, there may be no sequential equilibrium for some bargaining games. The non-existence problem of the equilibrium in pure strategies is critical since it may affect the applicability of alternating-offers protocol in realistic settings. Here we show one example which has no sequential equilibrium. The bargaining game with the following parameters: $T = 2$, $\iota(0) = \mathbf{b}$, $RP_s = 10$, $RP_1 = 100$, $RP_2 = 40$, $\omega_{\mathbf{b}_1}^0 = 0.6$, $\omega_{\mathbf{b}_2}^0 = 0.4$, $\delta_s = 0.9$, and $\delta_b = 0.8$. First consider that the buyer is applying the pooling accepting choice rule at time 0. With the initial belief, the seller will get a utility of $0.6(100 - 10)0.9^2 = 43.74$ if it offer RP_1 at time $t = 1$ and will get a utility of $(40 - 10)0.9^2 = 24.3$ if it offer RP_2 at time $t = 1$.

Table 3.1. Simulation parameters

<i>Parameters</i>	<i>Value range</i>
<i>Deadline (T)</i>	[2, 14]
<i>Number of buyer types ($\Delta_{\mathbf{b}}^0$)</i>	[2, 9]
<i>Reserve price of buyer (RP_i)</i>	[40, 100]
<i>Reserve price of seller (RP_s)</i>	[5, 20]
<i>Discounting factors ($\delta_{\mathbf{s}}, \delta_{\mathbf{b}}$)</i>	[0.5, 1]

Therefore, the optimal offer of the seller with the initial belief at $t = 1$ is offering 100. Accordingly, the optimal offer of all buyer types is $(100)_{\leftarrow \mathbf{s}} = 58.6$. Obviously, buyer \mathbf{b}_2 will deviate from it since the offering price 58.6 is higher than its reserve price 40. Consider the pooling rejecting choice rule in which all buyer types offer -1 and buyer \mathbf{b}_1 's equivalent offer is $(100)_{\leftarrow \mathbf{b}_1} = 100$. In this case, buyer \mathbf{b}_1 has an incentive to offer $x_{\mathbf{b}_1}^*(0) = 91$ which will be accepted by the seller since it will update its belief to $\{\mathbf{b}_1\}$ after receiving the offer other than 100. The final choice rule for the buyer is the separating choice rule in which buyer \mathbf{b}_1 offers $x_{\mathbf{b}_1}^*(0) = 91$ and buyer \mathbf{b}_2 offers -1 . However, buyer \mathbf{b}_1 has an incentive to offer -1 since at time $t = 1$ the seller will offer $\text{RP}_2 = 40$, which is better since $(40)_{\leftarrow \mathbf{b}_1} = 52 < x_{\mathbf{b}_1}^*(0) = 91$. Therefore, there is no sequential equilibrium for this bargaining game.

To evaluate the percentage of games with at least one pure strategy sequential equilibrium, we performed a series of experiments in a variety of test environments and the parameters are given in Table 6.2. In the experiments, the negotiation deadline is randomly selected from [2, 14], the number of buyer types is randomly selected from [2, 9] and the initial probability of each buyer type is set randomly. The reserve price of each buyer type is randomly selected from [40, 100] and the reserve price of the seller is randomly selected from [5, 20]. Therefore, the reserve price of each buyer is always higher than the reserve price of the seller and the two agents have a negotiation space. Agents' discounting factors model how agents' utilities decrease with time.

Table 3.2. Average number of sequential equilibria and percentage of games with sequential equilibria

T	$ \Delta_b^0 = 2$	$ \Delta_b^0 = 3$	$ \Delta_b^0 = 4$	$ \Delta_b^0 = 5$	$ \Delta_b^0 = 6$	$ \Delta_b^0 = 7$	$ \Delta_b^0 = 8$	$ \Delta_b^0 = 9$
2	1.64 (95.1%)	1.72 (95.4%)	1.78 (96.5%)	1.81 (96.4%)	1.84 (96.4%)	1.87 (97.3%)	1.89 (97.9%)	1.90 (97.8%)
3	1.99 (100%)	2.07 (100%)	2.17 (100%)	2.24 (100%)	2.30 (100%)	2.37 (100%)	2.39 (100%)	2.48 (99.9%)
4	2.37 (99.9%)	2.83 (100%)	3.23 (100%)	3.58 (100%)	3.87 (100%)	4.24 (100%)	4.38 (100%)	4.58 (100%)
5	2.49 (100%)	3.11 (100%)	3.67 (100%)	4.18 (100%)	4.70 (100%)	5.15 (100%)	5.53 (100%)	5.93 (100%)
6	2.87 (99.8%)	3.79 (100%)	4.69 (100%)	5.52 (100%)	6.34 (100%)	6.97 (100%)	7.66 (100%)	8.46 (100%)
7	3.06 (100%)	4.24 (100%)	5.25 (100%)	6.21 (100%)	7.16 (100%)	8.12 (100%)	9.16 (100%)	9.99 (100%)
8	3.54 (99.9%)	5.18 (100%)	6.59 (100%)	8.04 (100%)	9.67 (100%)	10.99 (100%)	12.55 (100%)	13.66 (100%)
9	3.82 (100%)	5.85 (100%)	7.38 (100%)	8.89 (100%)	10.62 (100%)	11.91 (100%)	13.95 (100%)	16.61 (100%)
10	4.55 (100%)	7.21 (100%)	9.92 (100%)	12.60 (100%)	15.48 (100%)	18.49 (100%)	20.71 (100%)	22.58 (100%)
11	5.04 (100%)	8.17 (100%)	11.11 (100%)	13.48 (100%)	17.22 (100%)	19.83 (100%)	25.15 (100%)	28.86 (100%)
12	6.36 (99.9%)	10.22 (100%)	16.54 (100%)	21.28 (100%)	27.57 (100%)	36.55 (100%)	38.05 (100%)	54.30 (100%)
13	6.89 (100%)	12.42 (100%)	16.79 (100%)	21.34 (100%)	28.57 (100%)	43.98 (100%)	48.35 (100%)	70.04 (100%)
14	8.90 (100%)	16.88 (100%)	26.10 (100%)	44.12 (100%)	70.63 (100%)	103.28 (100%)	125.86 (100%)	155.49 (100%)

When discounting factors are smaller than 0.5, agents' utilities drastically decrease with time. In order to make our setting more realistic, agents' discounting factors are randomly selected from $[0.5, 1)$. The agent making the first offer is randomly determined. The above setting represents a wide range of scenarios.

Experimental results show that there is at least one sequential equilibrium in $\sim 99.7\%$ of the bargaining games. Table 3.2 shows the average number of sequential equilibria (including both games with sequential equilibria and without sequential equilibria) and percentage of games with sequential equilibria in negotiation games with different deadlines and number of buyer types. For each combination of deadlines and number of buyer types, we randomly generated over 10000 scenarios and computed the average values. We can see that the average number of sequential equilibria or the percentage of games with sequential equilibria increases with the increase of deadlines. Similarly, the average number of sequential equilibria slightly increase with the number of buyer types, which corresponds to our intuitions. It can be observed from Table 3.2 that when the deadline is longer than 2 and the number of buyer types is more than 3, all scenarios have at least one sequential equilibrium. With more buyer types, there are more choice rules and reject update rules and poten-

tially, there will be more sequential equilibria. However, more buyer types introduce more stringent equilibrium existence conditions since an equilibrium requires that each buyer type has an incentive to choose a different strategy.

Our approach can find all sequential equilibria for each bargaining game. Based on the computation reduction techniques, we just need to find sequential equilibria for at most $|\Delta_{\mathbf{b}}^0|$ continuation games $\Gamma(t, \Delta_{\mathbf{b}})$ at time t where $\Delta_{\mathbf{b}} \subseteq \Delta_{\mathbf{b}}^0$. Let $\Psi(t, \Delta_{\mathbf{b}})$ be the maximum number of sequential equilibria for the continuation game $\Gamma(t, \Delta_{\mathbf{b}})$. If $t = T - 1$, it follows that $\Psi(t, \Delta_{\mathbf{b}}) = 1$. Otherwise, we have the following: 1) If $\iota(t) = \mathbf{b}$, $\Psi(t, \Delta_{\mathbf{b}}) = \mathcal{O}(|\Delta_{\mathbf{b}}| \Psi(t + 1, \Delta_{\mathbf{b}})^2)$ since buyer types can try different choice rules and for each choice rule, we need to consider all the equilibrium combinations of at most two continuation games. 2) If $\iota(t) = \mathbf{s}$, $\Psi(t, \Delta_{\mathbf{b}}) < \mathcal{O}(|\Delta_{\mathbf{b}}| (\Psi(t + 1, \Delta_{\mathbf{b}})^{|\Delta_{\mathbf{b}}|}))$ since we need to consider $|\Delta_{\mathbf{b}}| + 1$ reject update rules and for each reject update rule, we need to consider all the sequential equilibria for the corresponding continuation game. We can see that the number of sequential equilibria may exponentially increase with the number $|\Delta_{\mathbf{b}}^0|$ of buyer types and the deadline T and our algorithm has a high computational complexity. In our backward induction approach, the equilibrium calculation for a continuation game with certain belief may be conducted multiple times. To avoid the repetition of equilibrium calculation in our approach, one can store known equilibrium strategies for each continuation game with certain belief. If there is no calculation repetition, the computational complexity of our approach is not higher than any other complete algorithm.

We experimentally evaluated the running time to our algorithm using the setting specified in Table 6.2. All experiments run on a PC with a 2.16 Ghz Intel Pentium Dual processor and 2 GB of memory. Experimental results show that the algorithm's running time increases with the deadlines and the number of buyer types. Table 3.3 shows the average time of computing all the sequential equilibria in a bargaining game. The average running time is only about 12 seconds when the minimum deadline of the

Table 3.3. Average computation time (in seconds)

T	$ \Delta_{\mathbf{b}}^0 = 2$	$ \Delta_{\mathbf{b}}^0 = 3$	$ \Delta_{\mathbf{b}}^0 = 4$	$ \Delta_{\mathbf{b}}^0 = 5$	$ \Delta_{\mathbf{b}}^0 = 6$	$ \Delta_{\mathbf{b}}^0 = 7$	$ \Delta_{\mathbf{b}}^0 = 8$	$ \Delta_{\mathbf{b}}^0 = 9$
2	0.01	0.01	0.02	0.04	0.05	0.08	0.11	0.14
3	0.02	0.05	0.11	0.21	0.37	0.60	0.95	1.38
4	0.03	0.10	0.26	0.53	1.02	1.77	2.93	4.64
5	0.05	0.17	0.45	0.97	1.90	3.39	5.78	9.35
6	0.07	0.24	0.66	1.47	2.99	5.42	9.43	15.50
7	0.08	0.33	0.90	2.07	4.22	7.85	13.82	23.09
8	0.10	0.41	1.18	2.71	5.66	10.73	19.00	31.89
9	0.12	0.51	1.46	3.45	7.27	13.78	24.90	42.61
10	0.15	0.60	1.80	4.32	9.14	17.82	32.56	54.40
11	0.17	0.73	2.16	5.24	11.39	22.33	41.32	70.60
12	0.20	0.86	2.63	6.46	14.39	28.12	51.27	92.55
13	0.23	1.02	3.09	7.75	17.28	36.79	69.48	119.52
14	0.26	1.19	3.73	10.14	23.70	50.38	95.47	176.75

two agents is 14 and the number of buyer types is 9. We can find that the computation time increases drastically with the increase of both the negotiation deadline and the number of buyer types.

3.8 The Value of Equilibrium Strategies

This chapter provides a solution to compute pure strategy equilibria in bilateral bargaining games with one-sided uncertainty. Given that the chapter deals with a known negotiation setting, it would be very interesting to see how the pure strategy sequential equilibrium solution would fare against heuristic based strategies. An agent's equilibrium strategy is optimal given the other agent's equilibrium strategy. Therefore, if one agent adopts the equilibrium strategy, the other agent cannot gain a higher utility by switching to any other strategy (e.g., heuristic based strategies). Even though, it is still interesting to investigate how close the performance is when an agent uses different strategies, e.g., equilibrium strategy and a variety of heuristics. Furthermore, it would also be interesting to compare the social welfare (the total utility of all agents) when agents adopt different strategies.




Figure 3.4. The buyer's different rates of concession

3.8.1 Heuristic Based Strategies

While there have been a variety of heuristic based strategies in the literature, the most widely used strategy is the time dependent strategy (e.g., [12, 48, 78, 124, 127]). In a time dependent strategy, an agent makes concessions in response to the remaining negotiation time. Formally, an agent's offer at time t is given by:

$$IP_{\mathbf{a}} + (RP_{\mathbf{a}} - IP_{\mathbf{a}}) \left(\frac{t}{T-1} \right)^{\varepsilon}$$

where $RP_{\mathbf{a}}$ is agent \mathbf{a} 's reserve price, T is the minimum of two agents' deadline, and $IP_{\mathbf{a}}$ is agent \mathbf{a} 's optimal (or initial) price. An agent's optimal price is the agent's favorable price. In our bilateral bargaining game, an agent's optimal price is the reserve price of the other negotiation agent.

With infinitely many values of ε , there are infinitely many possible strategies in making concessions with respect to the remaining time. However, they can be classified into: 1) *Linear*: $\varepsilon = 1$, 2) *Conciliatory*: $0 < \varepsilon < 1$, and 3) *Conservative*: $\varepsilon > 1$.

ε reflects an agent's mental state about its eagerness for finishing the negotiation earlier [48, 78]. Figure 3.4 shows the buyer's different rates of concession.

3.8.2 Different Strategy Combination

In this section, we measure the value of equilibrium strategies against different time dependent strategies. For the experiment, we choose three representative dependent strategies (see Figure 3.4): 1) $\varepsilon = 1$: a liner concession making strategy; 2) $\varepsilon = 0.25$: a conciliatory concession making strategy; and 3) $\varepsilon = 4$: a conservative concession making strategy. When the seller is adopting a time dependent strategy, it will use the buyer's expected reserve price $\sum_{1 \leq i \leq n} \omega_{\mathbf{b}_i}^0 \text{RP}_i$ as its optimal price. Specifically, we consider the following 16 strategy combinations:

1. Both agents adopt the equilibrium strategies computed by using our algorithm. This strategy combination is called **EQ~EQ**.
2. The seller adopts the equilibrium strategy together with its belief system, but the buyer takes a time dependent strategy. There are totally 3 strategy combinations **0.25~EQ**, **1~EQ**, and **4~EQ**, where **0.25~EQ** indicates that the buyer is using a conciliatory time dependent strategy (i.e., $\varepsilon = 0.25$) but the seller is adopting the equilibrium strategy.
3. The buyer adopts the equilibrium strategy but the seller takes a time dependent strategy. In this case, the seller will not update its belief since it uses the time dependent strategy. However, the buyer will assume that the seller will update its belief using the belief system of its equilibrium strategy. **EQ~0.25**, **EQ~1**, and **EQ~4**, where **EQ~0.25** indicates that the buyer is adopting the equilibrium strategy but the seller is using a conciliatory time dependent strategy.
4. Both agents adopt the time dependent strategy. Therefore, there are 9 combinations: **0.25~0.25**, **0.25~1**, **0.25~4**, **1~0.25**, **1~1**, **1~4**, **4~0.25**, **4~1**, and




Figure 3.5. Buyer’s average utility while using different strategies

4~4 where 1~4 indicates that the buyer is adopting a time dependent strategy $\varepsilon = 1$ and the seller is adopting a time dependent strategy $\varepsilon = 4$.

3.8.3 Performance Measures and Results

We performed a series of experiments in a variety of test environments using the parameters from Table 6.2. Over 10^6 bargaining games were randomly generated that each had at least one sequential equilibrium. For each sequential equilibrium of a bargaining game, we tried different strategy combinations and for each strategy combination, we computed the equilibrium bargaining outcome for each buyer type. Based on the equilibrium bargaining outcome for each buyer type, we also computed the average utility of all buyer types and the seller. Given the average utilities of both the buyer and the seller in each sequential equilibrium for a bargaining game, we can compute the average utilities of both the buyer and the seller in the bargaining game. The main performance measure is the average utility of the buyer or the seller in each




Figure 3.6. Buyer’s average utility and the negotiation deadline

strategy combination. Another important performance measure is the social welfare (i.e, the total utility of the buyer and the seller) in each strategy combination.

Extensive stochastic simulations were carried out for all the combinations of variables in Table 6.2. For each combination, we randomly generated over 10000 experiments and tried all the strategy combinations and generated average performance measures. Even though extensive stochastic simulations were carried out for all the situations to compare the performance of different heuristic strategies and equilibrium strategies, due to space limitations, we only present the representative results in the rest of this section. For most results, we did not report the confidence intervals in our discussions since we found that the confidence interval for each average value was very tight around the value (e.g., the confidence intervals in Figure 3.5, Figure 3.8, and Figure 3.10).




Figure 3.7. Buyer’s average utility and agents’ discount factors

3.8.4 The Value of Choosing Equilibrium strategies

The first question we investigated is what is an agent’s advantage of choosing an equilibrium strategy over heuristic strategies given that the other agent is adopting the equilibrium strategy. Figure 3.5 shows the buyer’s average utility (as well as the 99% confidence interval) while using different strategies given that the seller is always using the equilibrium strategy. We can see that the equilibrium strategy (i.e., the **EQ~EQ** strategy combination) achieved a much higher utility than the other three heuristic strategies, i.e., strategy combinations **0.25~EQ**, **1~EQ**, and **4~EQ**. Furthermore, the utility of using a conciliatory strategy (e.g., $\varepsilon = 0.25$) is higher than that of using a conservative strategy (e.g., $\varepsilon = 4$). This result is mainly due to the existence of discount factor since an earlier agreement (even with the same price) can bring the buyer a higher utility.

Figure 3.6 shows how the negotiation deadline affects the buyer’s utility when the seller is using the equilibrium strategy. It can be observed that when both agents adopt the equilibrium strategy, the negotiation deadline almost has no influence on




Figure 3.8. Seller’s average utility while using different strategies

the buyer’s utility, which is similar to the strategy combination **0.25~EQ** where the buyer adopts a conciliatory concession making approach. However, when the buyer is adopting the linear or conservative strategy, its utility decreases with the increase of deadline. This result corresponds to the intuition that, when an agent uses a conservative (or linear) strategy, it will make an agreement later than using a conciliatory strategy. Correspondingly, it will gain a lower utility due the existence of the discount factor.

In addition to negotiation deadlines, agents’ discount factors also affect the buyer’s utility. Figure 3.7 shows how the buyer’s utility changes with the sum $\delta = \delta_b + \delta_s$ of both agents’ discount factors. Since δ_a is in the range of $[0.5, 1)$, it follows that $\delta \in [1, 2)$. It can be seen from Figure 3.7 that for different strategy combinations, the buyer’s utility increases with the increase of δ . Intuitively, with a higher discount factor, the buyer gains a higher utility even for the same negotiation negotiation result.




Figure 3.9. Seller’s average utility and agents’ discount factors

The other situation is that the buyer is adopting the equilibrium strategy but the seller is using the equilibrium strategy or different heuristic strategies. Figure 3.8 shows the seller’s average utility while using different strategies given that the buyer is always using the equilibrium strategy. It is easy to see that the equilibrium strategy (i.e., the **EQ~EQ** strategy combination) achieved a much higher utility than the other three heuristic strategies, i.e., strategy combinations **EQ~0.25**, **EQ~1**, and **EQ~4**. Similar to the results in Figure 3.5, the utility of using a conciliatory strategy is higher than that of using a conservative strategy. Figure 3.7 shows how the seller’s utility changes with the sum $\delta = \delta_b + \delta_s$ of both agents’ discount factors. We can see that for different strategy combinations, the seller’s utility increases with the increase of δ , which is again similar to the results in Figure 3.9.

3.8.5 Comparison of Social Welfare

We also compared the social welfare of different strategy combinations. From Figure 3.10 we can see that the social welfare of the strategy combination **EQ~EQ** is




Figure 3.10. Social welfare for different strategy combinations

much higher than that of all other strategy combinations, which corresponds to the results that agents achieved higher utilities while adopting equilibrium strategies. We can also see that the social welfare increases when agents' strategies are more conciliatory. When agents adopt conciliatory strategies, they will make large concessions at the beginning and it is more likely that they can make an agreement earlier, which results in higher social welfare.

Figure 3.11 shows that the social welfare of each strategy combination increases with the increase of the sum of both agents' discount factors, which corresponds to the results in Figure 3.7 and Figure 3.9. The main reason is that an agent's utility of a negotiation outcome may increase with its discount factor. Note that even when the sum of discount factor is close to 2, the strategy combination **EQ~EQ** still achieved a higher social welfare than other strategy combinations.

From Figure 3.12, we can see that the negotiation deadline has almost no effect on the social welfare when both agents adopt the equilibrium strategies. However, if one agent is using a heuristic strategy, the social welfare decreases with the deadline,




Figure 3.11. Social welfare and agents' discount factors

which is mainly due to the delay of agreement making time. Especially, when agents' strategies are more conservative, the advantage of the strategy combination $\mathbf{EQ}\sim\mathbf{EQ}$ increases with the increase of negotiation deadline.

3.9 Applications of the Approach

Our approach can be used to compute pure strategy sequential equilibria in other games with continuous strategy space. This section briefly discusses two potential application of our proposed approach.

3.9.1 Bilateral Negotiation with Uncertain Discount Factor

This chapter considers one-sided uncertainty regarding reserve prices and we assume complete knowledge about agents' discount factors. In other cases, one agent may have incomplete information about the other agent's discount factors. For example, we can assume that the buyer \mathbf{b} can be of finitely many types $\{\mathbf{b}_1, \dots, \mathbf{b}_n\}$ in which buyer \mathbf{b}_i has a discount factor $\delta_{\mathbf{b}}^i$. The initial belief of \mathbf{s} on \mathbf{b} is $\mu(0) = \langle \Delta_{\mathbf{b}}^0, P_{\mathbf{b}}^0 \rangle$




Figure 3.12. Social welfare and the negotiation deadline

where $\Delta_{\mathbf{b}}^0 = \{\mathbf{b}_1, \dots, \mathbf{b}_n\}$ and $P_{\mathbf{b}}^0 = \{\omega_{\mathbf{b}_1}^0, \dots, \omega_{\mathbf{b}_n}^0\}$ such that $\sum_i \omega_{\mathbf{b}_i}^0 = 1$. $\omega_{\mathbf{b}_i}^0$ is the priori probability that \mathbf{b} is of type \mathbf{b}_i . Let $x_{\mathbf{b}_i}^*(t)$ be any agent optimal offer at time t when \mathbf{b} is of type \mathbf{b}_i in this case. It follows that $x_{\mathbf{b}_i}^*(t) \leq x_{\mathbf{b}_j}^*(t)$ if $\delta_{\mathbf{b}}^i > \delta_{\mathbf{b}}^j$.

Lemma 8. $x_{\mathbf{b}_i}^*(t) \leq x_{\mathbf{b}_j}^*(t)$ if $\delta_{\mathbf{b}}^i > \delta_{\mathbf{b}}^j$.

Proof. Case 1 ($\iota(T) = \mathbf{s}$). It follows that $x_{\mathbf{b}_i}^*(T-1) = x_{\mathbf{b}_j}^*(T-1) = \text{RP}_{\mathbf{s}}$. Then $x_{\mathbf{b}_i}^*(T-2) = \text{RP}_{\mathbf{b}} - \delta_{\mathbf{b}}^i(\text{RP}_{\mathbf{b}} - x_{\mathbf{b}_i}^*(T-1)) < x_{\mathbf{b}_j}^*(T-2) = \text{RP}_{\mathbf{b}} - \delta_{\mathbf{b}}^j(\text{RP}_{\mathbf{b}} - x_{\mathbf{b}_i}^*(T-1))$. Similarly, we have $x_{\mathbf{b}_i}^*(T-3) = \text{RP}_{\mathbf{s}}(1 - \delta_{\mathbf{s}}) + \delta_{\mathbf{s}}x_{\mathbf{b}_i}^*(T-2)$ and $x_{\mathbf{b}_j}^*(T-3) = \text{RP}_{\mathbf{s}}(1 - \delta_{\mathbf{s}}) + \delta_{\mathbf{s}}x_{\mathbf{b}_j}^*(T-2)$. Thus we have $x_{\mathbf{b}_i}^*(T-3) < x_{\mathbf{b}_j}^*(T-3)$. Recursively, we have $x_{\mathbf{b}_i}^*(t) < x_{\mathbf{b}_j}^*(t)$ for $t < T-3$.

Case 2 ($\iota(T) = \mathbf{b}$). It follows that $x_{\mathbf{b}_i}^*(T-1) = x_{\mathbf{b}_j}^*(T-1) = \text{RP}_{\mathbf{b}}$. Then at time $T-2$, we have $x_{\mathbf{b}_i}^*(T-2) = \text{RP}_{\mathbf{s}}(1 - \delta_{\mathbf{s}}) + \delta_{\mathbf{s}}x_{\mathbf{b}_i}^*(T-1) = x_{\mathbf{b}_j}^*(T-2)$. Thus, $x_{\mathbf{b}_i}^*(T-2) > x_{\mathbf{b}_j}^*(T-2)$. As in case 1, we have $x_{\mathbf{b}_i}^*(t) < x_{\mathbf{b}_j}^*(t)$ for $t < T-2$. \square

In this setting, it is easy to see that the weakest type is the buyer type with the lowest discount factor. Accordingly, if the buyer acts off the equilibrium path, the

seller will update its belief to the buyer type with the lowest discount factor following the optimistic conjecture.

One can directly apply our approach to solve the bargaining game with uncertain discount factors. For the bargaining game with uncertain reserve prices, we use two techniques to reduce the number of choice rules and reject update rules that need to be considered. Fortunately, we can still only need to a small number of choice rules and reject update rules for the bargaining game with uncertain discount factors.

First we consider the reject update rule such that the seller will update its belief to $\Delta'_b \subseteq \Delta_b$ if the seller's offer is rejected by the buyer where Δ_b is the seller's belief while making the offer.

Theorem 9. *If there is a reject update rule with updated belief $\Delta'_b \subseteq \Delta_b$ such that $\delta_b^i > \delta_b^j$ for $\mathbf{b}_i \in \Delta_b \setminus \Delta'_b$ and $\mathbf{b}_j \in \Delta'_b$, agents' strategies are not sequentially rational.*

Proof. The result can be proved by contradiction. Assume that there is a sequential equilibrium with this reject update rule. Similar to Theorem 1, equilibrium existence conditions requires that $(e_{\mathbf{b}_j}^{t+1} | \Delta'_b)_{\leftarrow \mathbf{b}_j} \leq (e_{\mathbf{b}_i}^{t+1} | \Delta'_b)_{\leftarrow \mathbf{b}_i}$ where $e_{\mathbf{b}_j}^{t+1} | \Delta'_b$ is \mathbf{b}_j 's equivalent offer in the continuation game starting from $t + 1$ with belief Δ'_b . Since $\delta_b^i > \delta_b^j$, it follows that $(e_{\mathbf{b}_j}^{t+1} | \Delta'_b)_{\leftarrow \mathbf{b}_j} > (e_{\mathbf{b}_i}^{t+1} | \Delta'_b)_{\leftarrow \mathbf{b}_i}$. There is a contradiction. \square

Accordingly, we just need to consider no more than $|\Delta_b|$ reject update rules without sacrificing completeness of our approach. Similarly, we just need to consider a small number of choice rules due to the following theorem.

Theorem 10. *Assume that \mathbf{b} behaves in different ways at a continuation game with belief set Δ_b where $\Delta_b = \Delta_b^a \cup \Delta_b^r$ at time t . If there is a buyer type $\mathbf{b}_i \in \Delta_b^a$ and a buyer $\mathbf{b}_j \in \Delta_b^r$ such that $\delta_b^i > \delta_b^j$, there is no sequential equilibrium for this choice rule.*

Proof. Similar to Theorem 6, this result can be proved by contradiction. If there is a sequential equilibrium, the following two conditions are satisfied: 1) Buyer type

\mathbf{b}_i has no incentive to behave as \mathbf{b}_j , i.e., $(e_s^{t+1}|\Delta_{\mathbf{b}}^a)_{\leftarrow s} \leq \rho(\mathbf{b}_i, t, EBO(\mathbf{b}_j, \wp^r))$; and
 2) Buyer type \mathbf{b}_j has no incentive to behave as \mathbf{b}_i , i.e., $\rho(\mathbf{b}_j, t, EBO(\mathbf{b}_j, \wp^r)) \leq (e_s^{t+1}|\Delta_{\mathbf{b}}^a)_{\leftarrow s}$. Therefore, equilibrium existence requires that $\rho(\mathbf{b}_j, t, EBO(\mathbf{b}_j, \wp^r)) \leq \rho(\mathbf{b}_i, t, EBO(\mathbf{b}_j, \wp^r))$.

Assume that $EBO(\mathbf{b}_j, \wp^r) = (x, t')$ where $T \geq t' > t$. From the definition of equivalent offers, we have $(RP_{\mathbf{b}} - \rho(\mathbf{b}_j, t, EBO(\mathbf{b}_j, \wp^r))) \cdot (\delta_{\mathbf{b}}^j)^{t+1} = (RP_{\mathbf{b}} - x) \cdot \delta_{\mathbf{b}}^{t'}$, which can be rewritten as $\rho(\mathbf{b}_j, t, EBO(\mathbf{b}_j, \wp^r)) = RP_{\mathbf{b}} - (RP_{\mathbf{b}} - x) \cdot (\delta_{\mathbf{b}}^j)^{t'-t-1}$. Similarly, we have $\rho(\mathbf{b}_i, t, EBO(\mathbf{b}_j, \wp^r)) = RP_{\mathbf{b}} - (RP_{\mathbf{b}} - x) \cdot (\delta_{\mathbf{b}}^i)^{t'-t-1}$. Since $\delta_{\mathbf{b}}^i > \delta_{\mathbf{b}}^j$, it follows that $\rho(\mathbf{b}_i, t, EBO(\mathbf{b}_j, \wp^r)) < \rho(\mathbf{b}_j, t, EBO(\mathbf{b}_j, \wp^r))$ which contradicts with equilibrium existence conditions. \square

In summary, our approach can be directly applied to solve the bargaining game with one-sided uncertain discount factor and we can use the similar techniques to reduce computational cost.

3.9.2 Bilateral Multi-issue Negotiation with Uncertain Weights

Another potential application of our approach is bilateral multi-issue negotiation, which is more complex and challenging than a single-issue negotiation since agents need to make tradeoffs between multiple issues. The problem of bargaining efficiently over multiple issues with complete information has been addressed in [50, 51, 52]. Agents' equilibrium strategies can be easily computed by extending the backward induction method in Section 3.2. Specifically, the acting agent $\iota(t)$ at time t chooses its best offer (consisting of values for each negotiation issue) that is acceptable to the other agent. In presence of incomplete information, it is common to compute agents' sequential equilibrium strategies. There are different sources of uncertainty. For uncertainty about agents' reserve prices, discount factors or negotiation deadlines, the calculation of sequential equilibria in a multi-issue negotiation game is the same as

that in a single issue negotiation game. The only new source of uncertainty introduced by having multiple issues is the uncertainty regarding the weights of different issues.

In multi-issue negotiation, two agents are negotiating over multiple issues $1, \dots, l$. For each issue i , let $\text{RP}_{\mathbf{a}}^i$ be agent \mathbf{a} 's reserve price for the issue. A negotiation outcome can be represented as $o = \langle o_1, \dots, o_l \rangle$. An agent \mathbf{a} 's utility of a negotiation outcome o is defined as $U_{\mathbf{a}}(o) = \sum_{1 \leq i \leq l} w_{\mathbf{a}}^i U_{\mathbf{a}}(o_i)$ where $U_{\mathbf{a}}(o_i)$ is \mathbf{a} 's utility given the negotiation outcome o_i for issue i , which is the same as the utility function in single issue negotiation (see Section 3.2). In the cumulative utility function, $w_{\mathbf{a}}^i$ is agent \mathbf{a} 's weight for issue i . Let $w_{\mathbf{a}} = \langle w_{\mathbf{a}}^1, \dots, w_{\mathbf{a}}^l \rangle$ be agent \mathbf{a} 's weight vector. We consider the one-sided uncertainty about the buyer's weights of the issues and all other parameters are complete information. There are n possible weight vectors $\{w_{\mathbf{b}_1}, \dots, w_{\mathbf{b}_n}\}$ for the buyer and the probability of the buyer being the type $w_{\mathbf{b}_i}$ is $\omega_{\mathbf{b}_i}$. The probability distribution is common knowledge.

Fatima *et al.* [51, 52] present an algorithm to produce equilibrium strategies in multi-issue bargaining with uncertain weights but the strategies found by their algorithm are not necessarily sequentially rational given the designed system of beliefs as we discussed previously.⁵ Here we show a simple example (see Figure 3.13) where buyer \mathbf{b} and seller \mathbf{s} are negotiating over two issues 1 and 2 with the following parameters: $T = 5$, $\iota(0) = \mathbf{s}$, $\text{RP}_{\mathbf{s}}^1 = \text{RP}_{\mathbf{s}}^2 = 0$, $\text{RP}_{\mathbf{b}}^1 = \text{RP}_{\mathbf{b}}^2 = 90$, $\delta_{\mathbf{s}} = 0.5$, $\delta_{\mathbf{b}} = 0.8$, $w_{\mathbf{s}} = \langle 0.9, 0.1 \rangle$. There are 2 possible weight vectors $\{w_{\mathbf{b}_1} = \langle 0.4, 0.6 \rangle, w_{\mathbf{b}_2} = \langle 0.9, 0.1 \rangle\}$ for the buyer and the probability of the buyer being the type $w_{\mathbf{b}_1}$ and $w_{\mathbf{b}_2}$ are $\omega_{\mathbf{b}_1} = 0.9$ and $\omega_{\mathbf{b}_2} = 0.1$, respectively. At time $t = 4$, seller \mathbf{s} will offer buyer's reserve prices for both issues $\langle 90, 90 \rangle$, which gives the seller a utility of $0.9 \cdot (90 - 0) \cdot 0.5^5 + 0.1 \cdot (90 - 0) \cdot 0.5^5 = 2.8125$. At time $t = 3$, both buyer types'

⁵The multi-issue negotiation model here is slightly different from the multi-issue negotiation model in [51, 52] where two negotiation agents are splitting pies and the size of each pie shrinks over time due to the discount factors. In contrast, the utility of each agent shrinks over time. Our formulation of multi-issue negotiation has been widely used, e.g., [65, 50], just to name a few.




Figure 3.13. Failure of the approach in [52, 51] with $T = 5$, $\iota(0) = \mathbf{s}$, $RP_s^1 = RP_s^2 = 0$, $RP_b^1 = RP_b^2 = 90$, $\delta_s = 0.5$, $\delta_b = 0.8$, $w_s = \langle 0.9, 0.1 \rangle$, $\{w_{b_1} = \langle 0.4, 0.6 \rangle, w_{b_2} = \langle 0.9, 0.1 \rangle\}$, $\omega_{b_1} = 0.9$, and $\omega_{b_2} = 0.1$; agents' offers in complete information settings were also showed.

optimal offer that is acceptable to the seller is $\langle 50, 0 \rangle$. At time $t = 2$, seller's optimal offer that is acceptable to buyer type \mathbf{b}_1 is $\langle 85, 0 \rangle$ which can give seller an expected utility of 8.8875. Seller's optimal offer that is acceptable to buyer type \mathbf{b}_2 is $\langle 50, 90 \rangle$ which can give seller an expected utility of 3.20625. Therefore, seller's optimal offer at time $t = 2$ is $\langle 85, 0 \rangle$. In the same way, we can compute that both both buyer types' optimal offer at time $t = 1$ is $\langle 39.5, 0 \rangle$. Similarly, we can compute that seller's optimal offer at time $t = 0$ is $\langle 76.6, 0 \rangle$, which is only acceptable to the buyer type \mathbf{b}_1 . That is, if the buyer rejects the offer, the seller will update is belief to $\{\mathbf{b}_2\}$. However, buyer type \mathbf{b}_1 has an incentive to reject the offer and to make buyer type \mathbf{b}_2 complete information offer $\langle 30, 0 \rangle$ which can give \mathbf{b}_1 a higher utility than accepting seller's offer $\langle 76.6, 0 \rangle$.

We can apply our approach to solve the multi-issue bargaining game with uncertain issue weights in Figure 3.13. At time $t = 0$, we need to try different reject update rules and for each reject update rule, we first compute the sequential equilibria for its continuation game starting from time $t = 1$. To compute all the sequential equilibria for a continuation game starting from time $t = 1$, we need to consider different choice rules. While computing agents' equilibrium offers, one important optimization problem is computing one buyer type's (or a seller's) optimal offer which can give the seller (or a buyer type) certain utility. For instance, seller's expected utility is $y > 0$ and buyer type \mathbf{b}_j with weights $w_{\mathbf{b}_j} = \langle w_{\mathbf{b}_j}^1, \dots, w_{\mathbf{b}_j}^l \rangle$ is finding a package offer $x = \langle x^1, \dots, x^l \rangle$ at time t to maximize its utility. The optimization problem can be formulated as

$$\begin{aligned}
& \text{maximise} && \sum_{1 \leq i \leq l} w_{\mathbf{b}_j}^i \cdot (\text{RP}_{\mathbf{b}}^i - x^i) \cdot (\delta_{\mathbf{b}})^{t+1} \\
& \text{such that} && \sum_{1 \leq i \leq l} w_{\mathbf{s}}^i \cdot (x^i - \text{RP}_{\mathbf{s}}^i) \cdot (\delta_{\mathbf{s}})^{t+1} \geq y \\
& && \text{RP}_{\mathbf{b}}^i \leq x^i \leq \text{RP}_{\mathbf{b}}^i \text{ for } 1 \leq i \leq l
\end{aligned}$$

As implied in [51, 52], the optimal solution can be generated using a greedy approach by considering $w_{\mathbf{s}}^i/w_{\mathbf{b}_j}^i$ for $1 \leq i \leq l$ where $w_{\mathbf{b}_j}^i/w_{\mathbf{s}}^i$ is the utility that \mathbf{b}_j needs to give up in order increase \mathbf{s} 's utility by $(\delta_{\mathbf{b}})^{t+1}/(\delta_{\mathbf{s}})^{t+1}$. Thus, \mathbf{b}_j begins by making concessions to seller \mathbf{s} on the issue with the lowest $w_{\mathbf{b}_j}^i/w_{\mathbf{s}}^i$ value. Accordingly, the complexity of solving the optimization problem is polynomial.

The weakness of different buyer types at any continuation game can be computed by solving the bargaining game with complete information about weights of negotiation issues. Recall that that the weakest type gives the seller the highest utility in the complete information bargaining setting. Rather than only considering a small set of reject update rules and choice rules, here we may need to consider all reject

update rules and choice rules since the two computation reduction techniques are not necessarily valid here. However, this only increases the number of computations and does not affect the applicability of our approach. One sequential equilibrium for the multi-issue bargaining game in Figure 3.13 is that at time $t = 0$, seller \mathbf{s} makes an offer $\langle 69, 0 \rangle$ that is only acceptable to buyer type \mathbf{b}_1 . It follows that 1) offer $\langle 69, 0 \rangle$ is the seller’s optimal offer for this reject update rule; 2) buyer type \mathbf{b}_1 gains a utility of 49.92 and it has no incentive to reject the offer; and 3) by rejecting the offer, buyer type \mathbf{b}_2 will gain a utility of 40.32, which is higher than the utility 22.32 by accepting the offer $\langle 69, 0 \rangle$.

3.10 Summary

Studying rational agents’ strategic behavior is currently one of the most interesting issue in the field of automated negotiation. However, the bargaining theory literature lacks of general solutions for bargaining game with the presence of deadlines and incomplete information. In this chapter we go beyond state of the art by providing an algorithm that can find all sequential equilibria in incomplete information bargaining games with deadline constraints. Specifically, this chapter analyzes agents’ rational strategic behavior in alternating-offers bilateral bargaining with deadline constraints and one-sided uncertainty on reserve prices. Our approach computes sequential equilibrium employing a Bayesian extension of backward induction. To guarantee the completeness of our approach, we enumerate all choice rules and belief reject update rules. To guarantee the soundness of our approach, we construct equilibrium existence conditions along the backward induction process. Our approach can also be applied to other uncertainty settings, e.g., bilateral multi-issue negotiation with uncertain weight functions [51, 52], and bilateral bargaining with uncertain discount factors. We also compared the performance of the equilibrium strategies and representative heuristic based strategies. Empirical results show that agents with equilibrium strate-

gies achieved higher utilities than agents with heuristic based strategies. Furthermore, when both agents adopt the equilibrium strategies, the agents achieved much higher social welfare than that in all other strategy combinations.

Our study shows that there exists at least one sequential equilibrium in more than 99.7% of scenarios we have tried in which there are deadline constraints and incomplete information. There are two future research directions for this equilibrium nonexistence problem. On one hand, we can develop algorithm for finding mixed equilibrium strategies for bargaining scenarios in which there is no pure strategy equilibrium. On the other hand, we can slightly modify the alternating-offers protocol that would allow the existence of the equilibrium in pure strategies, e.g., the introduction of agents' strategic delay option [41].

While this chapter only considers one-sided uncertainty, we think our approach can be extended to handle two-sided uncertainty. Assume that the buyer is also uncertain about the seller's reserve price. When it is the buyer's turn to make an offer, rather than considering whether a buyer type's offer will be accepted by the seller, we need to consider the set of seller types that will accept the buyer's offer. That is, we need to combine choice rules and reject update rules. Similar to the buyer types, all seller types need to not only consider different reject update rules but also different choice rules. In addition to two-sided uncertainty regarding reserve prices, our algorithm can also handle two-sided uncertainty about discount factors.

One major motivation of the study of bargaining theory is designing successful bargaining agents in practical markets where there are more uncertainty and more agents. Although constraints, complexity, and uncertainty make it impractical to develop optimal negotiation strategies, our analysis can still give us some insights into the bargaining problems. Consider that a buyer is acquiring multiple resources in a dynamic market with multiple sellers. We can first use our approach to generate the strategy for each single seller and then use heuristics to combine the set of strategies

for all sellers to generate the overall negotiation strategy. The next chapter will look at concurrent negotiation between multiple buyers and sellers.

CHAPTER 4

EXTENDING ALTERNATING-OFFERS BARGAINING IN ONE-TO-MANY AND MANY-TO-MANY SETTINGS

The focus of this chapter is on analyzing agents' strategic behavior in one-to-many and many-to-many negotiations in which agents are negotiating with multiple trading partners and, at the same time, are facing competition from trading competitors. The subgame perfect equilibrium for complete information setting is presented and equilibrium properties, such as uniqueness, are discussed. We also analyze the reduction of computation to find sequential equilibria in one-to-many settings and many-to-many settings. Furthermore, we provide an algorithm to compute the sequential equilibrium in the incomplete information setting where there is uncertainty regarding the reserve price of an agent. This latter work will build on techniques developed in the previous chapter. The main goal of this chapter is to begin to understand which factors are affecting agents' bargaining position relative to others when each agent is negotiating with multiple trading partners simultaneously. This chapter is the first work to provide a game theoretical analysis of agents' strategic interactions in concurrent negotiations.

4.1 One-to-Many Alternating-Offers Negotiation

4.1.1 Negotiation Mechanism

In this section, we extend the alternating-offers protocol to capture the situation wherein there is one buyer agent \mathbf{b} and a set $\mathcal{S} = \{\mathbf{s}_1, \dots, \mathbf{s}_n\}$ of n seller agents such that: 1) the items sold by the sellers are the same, 2) all the sellers have exactly one item to sell, and 3) the buyer is interested in buying exactly one item.

Our mechanism extends the alternating-offers protocol allowing the buyer to carry on more simultaneous negotiations, each one with a different seller. As in [97, 125], a buyer synchronously negotiates with multiple sellers in discrete time. We use the term “negotiation thread” for the single bargaining between \mathbf{b} and a seller \mathbf{s}_i and we denote it by $\mathfrak{S}_{\mathbf{b},\mathbf{s}_i}$. Furthermore, we denote by $\iota(\mathfrak{S}_{\mathbf{b},\mathbf{s}_i}, t)$ the agent that acts at t in the negotiation thread $\mathfrak{S}_{\mathbf{b},\mathbf{s}_i}$. We assume that if $\iota(\mathfrak{S}_{\mathbf{b},\mathbf{s}_i}, t) = \mathbf{b}$ then $\iota(\mathfrak{S}_{\mathbf{b},\mathbf{s}_j}, t) = \mathbf{b}$ for all j . That is, \mathbf{b} simultaneously acts in all the negotiation threads. Therefore, if \mathbf{b} is proposing at time t , $\iota(t) = \mathbf{b}$. Otherwise, $\iota(t) = \mathcal{S}$.

We modify the alternating-offers mechanism by introducing an action *confirm* to avoid agents’ non-reasonable behaviors. In the following we show an example of non-reasonable behavior in absence of such action. The sellers’ action space is $A = \{offer[x], accept, exit, confirm\}$, whereas the buyer’s action space is the Cartesian product $\times_{i=1}^n A$. Legal actions for the buyer are all the pure strategies $\sigma_{\mathbf{b}} = \langle \sigma_{\mathbf{b},\mathbf{s}_1}, \dots, \sigma_{\mathbf{b},\mathbf{s}_n} \rangle$ such that: if $\sigma_{\mathbf{s}_i}(t-1) \neq accept$, then $\sigma_{\mathbf{b},\mathbf{s}_i}(t) \in \{offer[x], accept, exit\}$ except when $t = 0$, *accept* is not available, otherwise $\sigma_{\mathbf{b},\mathbf{s}_i}(t) \in \{confirm, exit\}$. Legal actions for the sellers are defined analogously: if $\sigma_{\mathbf{b},\mathbf{s}_i}(t-1) \neq accept$, then $\sigma_{\mathbf{s}_i}(t) \in \{offer[x], accept, exit\}$ except when $t = 0$, *accept* is not available, otherwise $\sigma_{\mathbf{s}_i}(t) \in \{confirm, exit\}$. The action *confirm* is allowed only after making the action *accept*.

The outcome of a single negotiation thread $\mathfrak{S}_{\mathbf{b},\mathbf{s}_i}$ is *NoAgreement* if either \mathbf{b} or \mathbf{s}_i made *exit*, whereas it is an agreement (x, t) if $\sigma_{\iota(\mathfrak{S}_{\mathbf{b},\mathbf{s}_i}, t)}(t) = confirm$, where x is such that $\sigma_{\iota(\mathfrak{S}_{\mathbf{b},\mathbf{s}_i}, t-2)}(t-2) = offer[x]$. Notice that, in absence of the action *confirm*, if \mathbf{b} makes offers to multiple sellers and all these accept, \mathbf{b} must buy multiple items. In presence of the action *confirm*, \mathbf{b} is in the position to choose only one contract. Summarily, in our mechanism the following process is needed for implementing an agreement: one agent proposes a price, the other agent accepts the offer, then the first

agent confirms the contract made by the second agent. Without loss of generality, we assume that each seller's deadline is no less than 2, i.e., $T_{s_i} \geq 2$.

The utility functions of the seller agents are exactly those defined in the previous section. However, we need to refine the utility function of \mathbf{b} . This is because \mathbf{b} can potentially buy more items, but is interested in only one item. We redefine \mathbf{b} 's utility as follows. If \mathbf{b} has reached more than one agreement, let (x_{first}, t_{first}) be the agreement such that, for any other agreement (x_j, t_j) , (1) $t_{first} \leq t_j$ and (2) $x_{first} \leq x_j$ if $t_{first} = t_j$. Let i_{first} be the seller involved in the agreement (x_{first}, t_{first}) . Agent \mathbf{b} 's utility is defined over the set of agreements it reached:

$$U_{\mathbf{b}}(\{(x_i, t_i)\}) = \begin{cases} (\text{RP}_{\mathbf{b}} - x_{first}) \cdot \delta_{\mathbf{b}}^{t_{first}} - \sum_{j \neq i_{first}} x_j & \text{if } t_{first} \leq T_{\mathbf{b}} \\ -\epsilon & \text{otherwise} \end{cases}$$

That is, \mathbf{b} receives a positive utility from the first agreement, whereas all the other agreements reduce \mathbf{b} 's utility. This will induce a rational buyer to reach at most one agreement.

4.1.2 Agents' Equilibrium Strategies

Let $\mathcal{S}_{=t}$ be the set of sellers whose deadline is t , i.e., $\mathcal{S}_{=t} = \{\mathbf{s}_i | T_{s_i} = t\}$. Let \mathcal{S}_t be the set of sellers which have no shorter deadline than t , i.e., $\mathcal{S}_t = \{\mathbf{s}_i | T_{s_i} \geq t\} = \cup_{t' \geq t} \mathcal{S}_{=t'}$. Without loss of generality, we assume that the sellers \mathcal{S}_t are ranked according to their reserve prices. We denote by \mathcal{S}_t^i ($\mathcal{S}_{=t}^i$) the seller with the i^{th} lowest reserve price in \mathcal{S}_t ($\mathcal{S}_{=t}$). Let $x_{\mathbf{b}, \mathbf{s}_i}^*(t)$ be \mathbf{b} 's optimal offer to \mathbf{s}_i at time t if $\iota(\mathfrak{S}_{\mathbf{b}, \mathbf{s}_i}, t) = \mathbf{b}$ and $x_{\mathbf{s}_i, \mathbf{b}}^*(t)$ be \mathbf{s}_i 's optimal offer to agent \mathbf{b} at time t if $\iota(\mathfrak{S}_{\mathbf{b}, \mathbf{s}_i}, t) = \mathbf{s}_i$.

The negotiation deadline for the negotiation thread between \mathbf{b} and \mathbf{s}_i is $T_{\mathbf{b}, \mathbf{s}_i} = \min(T_{\mathbf{b}}, T_{s_i})$. After $T_{\mathbf{b}, \mathbf{s}_i}$, at least one agent will have no interest in reaching agreements. Obviously, the negotiation deadline for \mathbf{b} is $T = \max_{\mathbf{s}_i \in \mathcal{S}} \{T_{\mathbf{b}, \mathbf{s}_i}\}$. We state the following lemma that allows us to reduce the complexity of the problem.

Lemma 11. *It is \mathbf{b} 's weakly dominant strategy to make the same offer to all the sellers in \mathcal{S}_{t+2} at each time t .*

Proof. At t we consider only \mathcal{S}_{t+2} since all the other sellers will not be interested in reaching agreements at $t+2$ and later. Consider the time point t wherein $\iota(t) = \mathbf{b}$. On the equilibrium path, at t agent \mathbf{b} will expect to reach exactly one agreement, say $(x_{\mathbf{b}}^*(t+2), t+2)$, with a specific seller, say \mathbf{s}^* . Obviously, \mathbf{s}^* is the seller that will accept the lowest offer. If \mathbf{b} makes offers higher than $x_{\mathbf{b}}^*(t)$ to the other sellers, then these sellers will not accept such offers and therefore \mathbf{b} cannot improve its utility. Analogously, if \mathbf{b} makes offers lower than $x_{\mathbf{b}}^*(t)$ to the other sellers, it cannot improve its utility. \square

According to Lemma 11 we can assume, without loss of generality, that $x_{\mathbf{b},\mathbf{s}_i}^*(t) = x_{\mathbf{b},\mathbf{s}_j}^*(t)$ for all $\mathbf{s}_i, \mathbf{s}_j$. For simplicity, we denote such offer by $x_{\mathbf{b}}^*(t)$. We state the following theorem.

Theorem 12. *In the one-to-many negotiation, the sequences of equilibrium offers $x_{\mathbf{b}}^*$ and $x_{\mathbf{s}_i}^*$ are:*

$$x_{\mathbf{b}}^*(t) = \begin{cases} \text{RP}_{\mathcal{S}_{t+2}^1} & t = T - 2 \text{ or } t = T_{\mathcal{S}_{t+2}^1} - 2 \\ \min\{(x_{\mathcal{S}_{t+2}^1}^*(t+1))_{\leftarrow \mathcal{S}_{t+2}^1}, \text{RP}_{\mathcal{S}_{t+2}^2}\} & t < T - 2 \text{ and } t \neq T_{\mathcal{S}_{t+2}^1} - 2 \end{cases},$$

$$x_{\mathbf{s}_i}^*(t) = \begin{cases} \max\{\text{RP}_{\mathbf{s}_i}, \text{RP}_{\mathcal{S}_T^2}\} & t = T - 2 \\ \max\{\text{RP}_{\mathbf{s}_i}, \min\{\text{RP}_{\mathcal{S}_{t+2}^2}, (x_{\mathbf{b}}^*(t+1))_{\leftarrow \mathbf{b}}\}\} & t < T - 2 \end{cases}.$$

Agent's equilibrium strategies are similar to those discussed in bilateral negotiation in Chapter 3, but $\sigma_{\mathbf{b},\mathbf{s}_i}$ prescribes that:

- \mathbf{b} accepts the offer x made by \mathbf{s}_i at t if: $x \leq (x_{\mathbf{b}}^*(t))_{\leftarrow \mathbf{b}}$ and x is the lowest received offer. If more than one seller has offered x , than \mathbf{b} accepts the offer made by the seller with the lowest reserve price;

- **b** confirms an accept of \mathbf{s}_i at t if: $\sigma_{\mathbf{b}}(t-2) = \text{offer}[x]$ with $x \leq (x_{\mathbf{b}}^*(t))_{\leftarrow 2[\mathbf{b}]}$ and, among all the sellers that have accepted $\sigma_{\mathbf{b}}(t-2)$, \mathbf{s}_i is the one with the lowest reserve price;

and $\sigma_{\mathbf{b}, \mathbf{s}_i}$ prescribes that:

- \mathbf{s}_i confirms the accept of **b** at t if: $\sigma_{\mathbf{s}_i}(t-2) = \text{offer}[x]$ with $x \geq \max\{(x_{\mathbf{s}_i}^*(t))_{\leftarrow 2[\mathbf{s}_i]}, \text{RP}_{\mathbf{s}_i}\}$.

Proof. First compute agents' optimal offers using backward induction. Let $x_{\mathcal{S}}^*(t) = \min_{\mathbf{s}_i \in \mathcal{S}_{t+2}} x_{\mathbf{s}_i}^*(t)$ be \mathcal{S} 's highest optimal offer at t . It follows that $x_{\mathbf{s}_i}^*(t) = \max\{\text{RP}_{\mathbf{s}_i}, x_{\mathcal{S}}^*(t)\}$. At time point T , the game for the buyer **b** rationally stops. The equilibrium outcome of every subgame starting from $t \geq T$ is *NoAgreement*. Therefore, at $t = T$ agent $\iota_{\mathfrak{S}_{\mathbf{b}, \mathbf{s}_i}}(T)$ would only confirm the best agreement proposed by agent $\iota_{\mathfrak{S}_{\mathbf{b}, \mathbf{s}_i}}(T-1)$. At time $t = T-1$, $\iota_{\mathfrak{S}_{\mathbf{b}, \mathbf{s}_i}}(T-1)$ will accept the best offer by agent $\iota_{\mathfrak{S}_{\mathbf{b}, \mathbf{s}_i}}(T-2)$, if $\iota_{\mathfrak{S}_{\mathbf{b}, \mathbf{s}_i}}(T-1)$ can get a utility not worse than *NoAgreement* by accepting the best offer. Note that at time $T-1$ and T , no agent will propose a price as it takes at least three time points to implement a final contract.

Assume that $\iota_{\mathfrak{S}_{\mathbf{b}, \mathbf{s}_i}}(t) = \mathbf{b}$. If $t = T-2$ or $t = T_{\mathcal{S}_{t+2}^1} - 2$, **b**'s optimal price is $\text{RP}_{\mathcal{S}_{t+2}^1}$ and seller \mathcal{S}_{t+2}^1 will accept it as its deadline is approaching. At $t < T-2$, $\min_{\mathbf{s}_i \in \mathcal{S}_{t+3}} ((x_{\mathbf{s}_i}^*(t+1))_{\leftarrow \mathbf{s}_i})$ is surely acceptable to some sellers in \mathcal{S}_{t+3} . We also need to consider sellers $\mathcal{S}_{t+2} - \mathcal{S}_{t+3}$ with deadline $t+2$, who are willing to accept any offer which is no less than their reserve prices. Therefore, **b**'s optimal offer at time t is

$$x_{\mathbf{b}}^*(t) = \min\left\{ \min_{\mathbf{s}_i \in \mathcal{S}_{t+3}} ((x_{\mathbf{s}_i}^*(t+1))_{\leftarrow \mathbf{s}_i}), \min_{\mathbf{s}_i \in \mathcal{S}_{t+2} - \mathcal{S}_{t+3}} \text{RP}_{\mathbf{s}_i} \right\} \quad (4.1)$$

It is easy to see that $x_{\mathcal{S}_{t+3}^1}^*(t+1) \leq x_{\mathcal{S}_{t+3}^2}^*(t+1) = \text{RP}_{\mathcal{S}_{t+3}^2}^*(t+1)$. It follows that $\min_{\mathbf{s}_i \in \mathcal{S}_{t+3}} ((x_{\mathbf{s}_i}^*(t+1))_{\leftarrow \mathbf{s}_i}) = (x_{\mathcal{S}_{t+3}^1}^*(t+1))_{\leftarrow \mathcal{S}_{t+3}^1}$. As $t \neq T_{\mathcal{S}_{t+2}^1} - 2$, equation (4.1) can be rewritten as $\min\{(x_{\mathcal{S}_{t+2}^1}^*(t+1))_{\leftarrow \mathcal{S}_{t+2}^1}, \text{RP}_{\mathcal{S}_{t+2}^2}\}$. Therefore, $x_{\mathbf{b}}^*(t) = \min\{(x_{\mathcal{S}_{t+2}^1}^*(t+1))_{\leftarrow \mathcal{S}_{t+2}^1}, \text{RP}_{\mathcal{S}_{t+2}^2}\}$ if $t < T-2$ and $t \neq T_{\mathcal{S}_{t+2}^1} - 2$.

Assume that $\iota_{\mathfrak{S}_{\mathbf{b}, \mathbf{s}_i}}(t) = \mathbf{s}_i$. At time $t = T - 2$, the acceptable offer to buyer \mathbf{b} is $\text{RP}_{\mathcal{S}_T^2}$ as all sellers in \mathcal{S}_T^2 compete with each other to get a contract. Thus, \mathbf{s}_i 's optimal offer is $\max\{\text{RP}_{\mathbf{s}_i}, \text{RP}_{\mathcal{S}_T^2}\}$. At time $t < T - 2$, the acceptable offer to buyer \mathbf{b} is $(x_{\mathbf{b}}^*(t + 1))_{\leftarrow \mathbf{b}}$. However, \mathbf{s}_i needs to consider the competition among sellers then \mathbf{s}_i 's winning price should be no higher than $\text{RP}_{\mathcal{S}_{t+2}^2}$. Then \mathbf{s}_i 's optimal offer is $\max\{\text{RP}_{\mathbf{s}_i}, \min\{\text{RP}_{\mathcal{S}_{t+2}^2}, (x_{\mathbf{b}}^*(t + 1))_{\leftarrow \mathbf{b}}\}\}$.

Finally, agents' optimal actions can be easily defined on the basis of $x_{\mathbf{b}}^*(t)$ and $x_{\mathbf{s}_i}^*(t)$. When an agent decides to make an offer, it always proposes its optimal offer $(x_{\mathbf{b}}^*(t)$ or $x_{\mathbf{s}_i}^*(t))$. Buyer \mathbf{b} will accept an offer $\sigma_{\mathbf{s}_i}(t - 1)$ if $\sigma_{\mathbf{s}_i}(t - 1) \leq (x_{\mathbf{b}}^*(t))_{\leftarrow \mathbf{b}}$ and $\sigma_{\mathbf{s}_i}(t - 1)$ is no higher than other sellers' offers at time $t - 1$. It is possible that several sellers propose a same acceptable offer. The tie can be broken by choosing the lowest offer from the seller with the lowest reserve price (note that we assume that sellers have different reserve prices). If at time $t - 1$, seller \mathbf{s}_i agrees with \mathbf{b} 's offer $\sigma_{\mathbf{b}, \mathbf{s}_i}(t - 2)$, \mathbf{b} will confirm the agreement if $\sigma_{\mathbf{b}, \mathbf{s}_i}(t - 2) \leq \sigma_{\mathbf{b}, \mathbf{s}_j}(t - 2)$ if \mathbf{s}_j also agrees with \mathbf{b} 's offer at time $t - 1$. Again, there could be more than one agreement with the same lowest price. To make sure that \mathbf{b} only makes one final agreement, \mathbf{b} confirms the agreement from the seller with the lowest reserve price. The optimal actions of all the sellers can be defined analogously. For simplicity, we consider just agents' strategies on the equilibrium path. \square

The computational complexity of the backward induction is $\mathcal{O}(nT)$ as the backward induction will go through all the time points and at each time point, each agent has at most three possible optimal actions. The equilibrium agreement is reached at $t = 2$ between \mathbf{b} and \mathcal{S}_2^1 and it is $(x_{\mathbf{b}}^*(0), 2)$ if $\iota(0) = \mathbf{b}$ and $(x_{\mathcal{S}_2^1}^*(0), 2)$ otherwise. It can be easily observed that $\text{RP}_{\mathcal{S}_2^1} \leq x_{\mathbf{b}}^*(0), x_{\mathcal{S}_2^1}^*(0) \leq \text{RP}_{\mathcal{S}_2^2}$. The result about agreement price is intuitive in the following sense: obviously, the agreement price cannot be lower than each seller's reserve price. But it also cannot be higher than the second lowest price as, if so, there is at least another seller who is willing to sell for less and

make an agreement with the buyer. Therefore, market competition guarantees that the buyer can make an agreement by paying no more than $\text{RP}_{\mathcal{S}_2^2}$. The lower bound of agreement is due to the proposing ordering and agents' deadlines. For example, if $T = 2$ and the buyer proposes at time $t = 0$, the buyer will propose $\text{RP}_{\mathcal{S}_2^1}$ and the agent \mathcal{S}_2^1 will accept the offer at time $t = 1$. We can see that the market competition plays an important role in affecting negotiation results. The buyer can make an agreement with price at most $\text{RP}_{\mathcal{S}_2^2}$. With more sellers, the buyer can get better (at least not worse) negotiation result.

Let us remark an observation. Consider the situation wherein $\iota(0) = \mathcal{S}$ and $x_{\mathcal{S}_2^1}^* = \text{RP}_{\mathcal{S}_2^2}$. Although both \mathcal{S}_2^1 and \mathcal{S}_2^2 have the same equilibrium offer, i.e., $\text{RP}_{\mathcal{S}_2^2}$, the equilibrium strategy of \mathbf{b} prescribes that \mathbf{b} must accept only the offer made by \mathcal{S}_2^1 . In the case \mathbf{b} accepts the offer by \mathcal{S}_2^2 or randomizes over accepting those offers, \mathcal{S}_2^1 's optimal action at $t = 0$ does not exist, being $\lim_{\varepsilon \rightarrow 0} (\mathcal{S}_2^2 - \varepsilon)$ with $\varepsilon \neq 0$. We can state the following theorem which is a direct consequence of the above observation and of the equilibrium uniqueness in bilateral alternating-offers.

Theorem 13. *Agents' strategies on the equilibrium path are unique except when $\text{RP}_{\mathcal{S}_2^1} = \text{RP}_{\mathbf{s}_i}$ for more than one i .*

Notice that, when the reserve price of more sellers is equal to $\text{RP}_{\mathcal{S}_2^1}$, all these sellers will offer their reserve price and \mathbf{b} can accept any single offer among these. However, it can be easily observed that all the equilibria are equivalent in terms of agents' payoffs, \mathbf{b} receiving the same utility in all the equilibria. As we assume that agents have different reserve prices, the equilibrium is unique.

Figure 4.1 shows an example of backward induction construction with $\text{RP}_{\mathbf{b}} = 1$, $\text{RP}_{\mathbf{s}_1} = 0$, $\text{RP}_{\mathbf{s}_2} = 0.2$, $\delta_{\mathbf{b}} = 0.8$, $\delta_{\mathbf{s}_1} = 0.7$, $\delta_{\mathbf{s}_2} = 0.8$, $T_{\mathbf{b}} = 10$, $T_{\mathbf{s}_1} = 11$, $T_{\mathbf{s}_2} = 7$. We report in the figure for any time point t the optimal offer $x_{\mathbf{a}}^*(t)$ that $\iota(t)$ can make; the dashed lines are sellers' optimal offers if there is only one seller. The time point from which we can apply the backward induction method is $T = 10$




Figure 4.1. Backward induction construction with $RP_b = 1$, $RP_{s_1} = 0$, $RP_{s_2} = 0.2$, $\delta_b = 0.8$, $\delta_{s_1} = 0.7$, $\delta_{s_2} = 0.8$, $T_b = 10$, $T_{s_1} = 11$, $T_{s_2} = 7$; at each time point t the optimal offer $x_a^*(t)$ that $\iota(t)$ can make is marked; the dashed lines are sellers' optimal offer if there is only one seller.

at which \mathbf{b} will confirm the agreement made at $t = 9$. At $t = 9$ agent s_1 will accept any offer equal to or higher than its reserve price $RP_{s_1} = 0$. The optimal offer $x_b^*(8)$ of \mathbf{b} at $t = 8$ is thus $RP_{s_1} = 0$. s_1 's optimal offer $x_{s_1}^*(7)$ at $t = 7$ is $(x_b^*(8))_{\leftarrow b} = RP_b - (RP_b - x_b^*(8))\delta_b = 0.2$. \mathbf{b} 's optimal offer at time $t = 6$ is then $x_b^*(6) = (x_{s_1}^*(7))_{\leftarrow s_1} = 0.14$. At time $t = 5$, another seller s_2 can make an offer (note that $t = 5$ is the last time s_2 can make an offer as it needs another two rounds to accept and confirm an agreement). s_1 and s_2 will compete with each other and their optimal offers aren't $(x_b^*(6))_{\leftarrow b} = 0.312$ as one seller has an incentive to choose a lower price if the other seller choose $(x_b^*(6))_{\leftarrow b} = 0.312$. The equilibrium optimal

price for the two sellers is $x_{s_1}^*(5) = x_{s_2}^*(5) = \text{RP}_{\mathcal{S}_{t=5+2}^2} = \text{RP}_{s_2} = 0.2$. The process continues to the initial time point $t = 0$ where \mathbf{b} 's optimal offer is $x_{\mathbf{b}}^*(0) = 0.14$.

There are some other mechanisms which can be used to implement contracts between buyer \mathbf{b} and sellers \mathcal{S} . Here we compare our model with the following mechanisms:

- *Bilateral bargaining without outside option*: Rubinstein's bilateral bargaining does not offer any mechanism to capture competition between sellers. In order to compare outcomes from bilateral bargaining with respect to outcomes from our mechanism, suppose that \mathbf{b} is able to choose the seller with which to negotiate. In our mechanism the buyer \mathbf{b} gains as in bilateral bargaining without outside option when the sequence of optimal offers $x^*(t)$ in the bilateral negotiation between \mathbf{b} and \mathcal{S}_2^1 is such that $x_i^*(t) \leq \text{RP}_{\mathcal{S}_2^2}$, otherwise the buyer \mathbf{b} gains more in our mechanism.
- *Bilateral bargaining with outside option*: In our mechanism the buyer gains no less than in bilateral bargaining with outside option in which an agent can leave the bilateral negotiation it is currently carrying on and negotiate with a different opponent [24]. We report an example. Consider the situation where there are two sellers with the same reservation price $\text{RP}_{\mathbf{s}}$ and any deadline no smaller than 2. In bilateral bargaining with outside option the agreement price is strictly larger than $\text{RP}_{\mathbf{s}}$, instead in our protocol the agreement price is exactly $\text{RP}_{\mathbf{s}}$.
- *VCG auction*: Since VCG auction does not take into account any temporal issues (no deadline and no discount factor), we limit our comparison to the agreement price. In VCG mechanism the agreement price is exactly $\text{RP}_{\mathcal{S}_2^2}$, whereas in our bargaining model the buyer's agreement price falls between $[\text{RP}_{\mathcal{S}_2^1}, \text{RP}_{\mathcal{S}_2^2}]$.

That is, the buyer achieved higher utility within our model which is also efficient.

4.1.3 Equilibrium Outcome Computation and Uncertain Information

We initially focus on the computation of the equilibrium outcome with complete information. Although agents' equilibrium strategies depend on the values of the parameters of all the agents, for a large subset of the space of the parameters the equilibrium outcome depends on the values of a narrow number of parameters. We have the following theorem.

Theorem 14. *When 1) $T_{\mathcal{S}_2^2} > 2$ if $\iota(0) = \mathbf{b}$ and 2) $(\text{RP}_{\mathbf{s}})_{\leftarrow \mathcal{S}_2^1 \mathbf{b}} \geq \text{RP}_{\mathbf{s}}$ for any seller $\mathbf{s} \in \mathcal{S}$, the equilibrium outcome depends only on the parameters of \mathbf{b} (i.e., $\text{RP}_{\mathbf{b}}$, $\delta_{\mathbf{b}}$, $T_{\mathbf{b}}$), \mathcal{S}_2^1 (i.e., $\text{RP}_{\mathcal{S}_2^1}$, $\delta_{\mathcal{S}_2^1}$, $T_{\mathcal{S}_2^1}$), and on the reserve price $\text{RP}_{\mathcal{S}_2^2}$ of \mathcal{S}_2^2 . In these situations the equilibrium outcome can be produced as follows:*

1. *finding the sequence of the optimal offers (say $y(t)$) under the assumption that \mathcal{S}_2^1 is the unique seller, and*
2. *assigning $x_{\mathbf{b}}^*(0) = \min\{y(0), (\text{RP}_{\mathcal{S}_2^2})_{\leftarrow \mathcal{S}_2^1}\}$ if $\iota(0) = \mathbf{b}$ and assigning $x_{\mathcal{S}}^*(0) = \min\{y(0), \text{RP}_{\mathcal{S}_2^2}\}$ if $\iota(0) = \mathcal{S}$.*

Proof. Case 1 ($\iota(\min\{T_{\mathcal{S}_2^1}, T_{\mathbf{b}}\}) = \mathbf{b}$). Let $t' + 2 = \min\{T_{\mathcal{S}_2^1}, T_{\mathbf{b}}\}$. It's easy to see that $x_{\mathbf{b}}^*(t') = \text{RP}_{\mathcal{S}_2^1} = y(t')$. Then we have $x_{\mathcal{S}}^*(t' - 1) = \min\{(\text{RP}_{\mathcal{S}_2^1})_{\leftarrow \mathbf{b}}, \text{RP}_{\mathcal{S}_{t'+1}^2}\} = \min\{y(t' - 1), \text{RP}_{\mathcal{S}_{t'+1}^2}\}$.¹ At time $t' - 2$, we have $x_{\mathbf{b}}^*(t' - 2) = \min\{(\text{RP}_{\mathcal{S}_2^1})_{\leftarrow \mathbf{b} \mathcal{S}_2^1}, (\text{RP}_{\mathcal{S}_{t'+1}^2})_{\leftarrow \mathcal{S}_2^1}, \text{RP}_{\mathcal{S}_{=t'}^1}\} = \min\{y(t' - 2), (\text{RP}_{\mathcal{S}_{t'+1}^2})_{\leftarrow \mathcal{S}_2^1}, \text{RP}_{\mathcal{S}_{=t'}^1}\}$. At time $t' - 3$, we have $x_{\mathcal{S}}^*(t' - 3) = \min\{(y(t' - 2))_{\leftarrow \mathbf{b}}, (\text{RP}_{\mathcal{S}_{t'+1}^2})_{\leftarrow \mathcal{S}_2^1 \mathbf{b}}, (\text{RP}_{\mathcal{S}_{=t'}^1})_{\leftarrow \mathbf{b}}, \text{RP}_{\mathcal{S}_{t'-1}^2}\}$. It's obvious that $(\text{RP}_{\mathcal{S}_{=t'}^1})_{\leftarrow \mathbf{b}} \geq (\text{RP}_{\mathcal{S}_{=t'}^1}) \geq \text{RP}_{\mathcal{S}_{t'-1}^2}$. In addition, as we assume that $(\text{RP}_{\mathbf{s}})_{\leftarrow \mathcal{S}_2^1 \mathbf{b}} \geq \text{RP}_{\mathbf{s}}$, it follows that $(\text{RP}_{\mathcal{S}_{t'+1}^2})_{\leftarrow \mathcal{S}_2^1 \mathbf{b}} \geq \text{RP}_{\mathcal{S}_{t'+1}^2} \geq \text{RP}_{\mathcal{S}_{t'-1}^2}$. Then we have $x_{\mathcal{S}}^*(t' - 3) =$

¹For convenience, $\text{RP}_{\mathcal{S}_{t'+1}^2} = \infty$ if $|\mathcal{S}_{t'+1}| < 2$.

$\min\{y(t' - 3), \text{RP}_{\mathcal{S}_{t'-1}^2}\}$. Following this procedure, we have 1) if $\iota(0) = \mathcal{S}$, $x_{\mathcal{S}_2^1}^*(0) = \min\{y(0), \text{RP}_{\mathcal{S}_2^2}\}$; 2) if $\iota(0) = \mathbf{b}$, $x_{\mathbf{b}}^*(0) = \min\{y(0), (\text{RP}_{\mathcal{S}_2^2})_{\leftarrow \mathcal{S}_2^1}\}$ as $(\text{RP}_{\mathcal{S}_3^2})_{\leftarrow \mathcal{S}_2^1} = (\text{RP}_{\mathcal{S}_2^2})_{\leftarrow \mathcal{S}_2^1} \leq (\text{RP}_{\mathcal{S}_2^2}) \leq \text{RP}_{\mathcal{S}_2^3} \leq \text{RP}_{\mathcal{S}_{=2}^1}$ given that $T_{\mathcal{S}_2^2} > 2$.

Case 2 ($\iota(\min\{T_{\mathcal{S}_2^1}, T_{\mathbf{b}}\}) = \mathcal{S}$). Let $t' + 2 = \min\{T_{\mathcal{S}_2^1}, T_{\mathbf{b}}\}$. At time t' , there are two situations: 1) $|\mathcal{S}_{t'+2}| < 2$, which implies that $x_{\mathcal{S}}^*(t') = \text{RP}_{\mathbf{b}} = y(t')$; 2) Otherwise, $x_{\mathcal{S}}^*(t') = \text{RP}_{\mathcal{S}_{t'+2}^2}$. Therefore, $x_{\mathcal{S}}^*(t') = \min\{y(t'), \text{RP}_{\mathcal{S}_{t'+2}^2}\}$. At time $t' - 1$, it follows that $x_{\mathbf{b}}^*(t' - 1) = \min\{y(t' - 1), (\text{RP}_{\mathcal{S}_{t'+2}^2})_{\leftarrow \mathcal{S}_2^1}, \text{RP}_{\mathcal{S}_{=t'+1}^1}\}$. Then at time $t' - 2$, we have $x_{\mathcal{S}}^*(t' - 2) = \min\{y(t' - 2), (\text{RP}_{\mathcal{S}_{t'+2}^2})_{\leftarrow \mathcal{S}_2^1 \mathbf{b}}, (\text{RP}_{\mathcal{S}_{=t'+1}^1})_{\leftarrow \mathbf{b}}, \text{RP}_{\mathcal{S}_{=t'}^1}, \text{RP}_{\mathcal{S}_t^2}\}$. It easy to see that $\text{RP}_{\mathcal{S}_{=t'}^1} \geq \text{RP}_{\mathcal{S}_t^2}$. It's obvious that $(\text{RP}_{\mathcal{S}_{=t'+1}^1})_{\leftarrow \mathbf{b}} \geq (\text{RP}_{\mathcal{S}_{=t'+1}^1}) \geq \text{RP}_{\mathcal{S}_t^2}$. As we assume that $(\text{RP}_{\mathbf{s}})_{\leftarrow \mathcal{S}_2^1 \mathbf{b}} \geq \text{RP}_{\mathbf{s}}$, it follows that $(\text{RP}_{\mathcal{S}_{t'+2}^2})_{\leftarrow \mathcal{S}_2^1 \mathbf{b}} \geq \text{RP}_{\mathcal{S}_{t'+2}^2} \geq \text{RP}_{\mathcal{S}_t^2}$. Then we have $x_{\mathcal{S}}^*(t' - 2) = \min\{y(t' - 2), \text{RP}_{\mathcal{S}_t^2}\}$. Following this procedure, we have 1) if $\iota(0) = \mathcal{S}$, $x_{\mathcal{S}_2^1}^*(0) = \min\{y(0), \text{RP}_{\mathcal{S}_2^2}\}$; 2) if $\iota(0) = \mathbf{b}$, $x_{\mathbf{b}}^*(0) = \min\{y(0), (\text{RP}_{\mathcal{S}_2^2})_{\leftarrow \mathcal{S}_2^1}\}$ given that $T_{\mathcal{S}_2^2} > 2$. \square

This is to say that the equilibrium outcome does not depend on the values of $\delta_{\mathcal{S}_2^2}$, $T_{\mathcal{S}_2^2}$, and on the parameters of all the other sellers. This is of paramount importance since complex settings with a high degree of uncertainty can be easily solved when 1) $T_{\mathcal{S}_2^2} > 2$ if $\iota(0) = \mathbf{b}$ and 2) $(\text{RP}_{\mathbf{s}})_{\leftarrow \mathcal{S}_2^1 \mathbf{b}} \geq \text{RP}_{\mathbf{s}}$ for any seller $\mathbf{s} \in \mathcal{S}$. Indeed, the above algorithm produces the equilibrium outcome even when $\delta_{\mathcal{S}_2^i}$ with $i > 1$, $T_{\mathcal{S}_2^i}$ with $i > 1$, and $\text{RP}_{\mathcal{S}_2^i}$ with $i > 2$ are uncertain. We can write the condition $(\text{RP}_{\mathcal{S}_2^2})_{\leftarrow \mathcal{S}_2^1 \mathbf{b}} \geq \text{RP}_{\mathcal{S}_2^2}$ as

$$(\text{RP}_{\mathbf{b}} - \text{RP}_{\mathcal{S}_2^1}) \geq (\text{RP}_{\mathcal{S}_2^2} - \text{RP}_{\mathcal{S}_2^1}) \frac{1 - \delta_{\mathbf{b}} \delta_{\mathcal{S}_2^1}}{1 - \delta_{\mathbf{b}}}.$$

It can be easily observed that, in common real-world settings where $\text{RP}_{\mathbf{b}} \gg \text{RP}_{\mathcal{S}_2^2}$ and $\delta_{\mathcal{S}_2^1}$ is close to 1, the above condition is satisfied.

Now, we focus on the uncertainty over \mathbf{b} 's and \mathcal{S}_2^1 's parameters. The values of these parameters affect the equilibrium outcome and therefore in presence of uncertainty over them we need to compute agents' equilibrium strategies to derive the equilibrium outcome. Currently, the literature provides algorithms to compute agents' equilibrium

strategies only in bilateral settings without outside option with one-sided uncertainty over deadlines [61]. We recall that, since the number of available actions is infinite, no algorithms such as Lemke-Howson [135] can be employed to compute a sequential equilibrium.

When $\text{RP}_{\mathcal{S}_2^2} \leq (\text{RP}_{\mathcal{S}_2^2})_{\leftarrow \mathcal{S}_2^1 \mathbf{b}}$, the algorithm presented in [61] can be easily extended to capture uncertainty in one-to-many bargaining. More precisely, we have that:

- when $T_{\mathbf{b}}$ is uncertain, whereas $T_{\mathcal{S}_2^1}$ is certain, then agents' equilibrium strategies can be produced by employing the algorithm presented in [61] where the buyer is \mathbf{b} and the seller is \mathcal{S}_2^1 and upper bounding the optimal offers to $\text{RP}_{\mathcal{S}_2^2}$ if $\iota(0) = \mathbf{b}$ and to $(\text{RP}_{\mathcal{S}_2^2})_{\leftarrow \mathcal{S}_2^1}$ if $\iota(t) = \mathcal{S}$;
- when $T_{\mathcal{S}_2^1}$ is uncertain, whereas $T_{\mathbf{b}}$ is certain, then agents' equilibrium strategies can be computed.

Settings with a higher degree of uncertainty, such as when both $T_{\mathbf{b}}$ and $T_{\mathcal{S}_2^1}$ are uncertain, need further exploration.

The results discussed above show that the analytical complexity of one-to-many bargaining is drastically less complicated than that of bilateral bargaining with outside option. This allows one to drastically reduce the search space and makes the computation easy. Therefore, one-to-many bargaining seems more appropriate for real-world settings when computational issues should be considered.

4.2 Many-to-Many Alternating-Offers Negotiation

4.2.1 Negotiation Mechanism

In this section, we propose a bargaining model for many-to-many negotiation where m buyer agents $\mathcal{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_m\}$ negotiate n seller agents $\mathcal{S} = \{\mathbf{s}_1, \dots, \mathbf{s}_n\}$. In this case, both buyers and sellers face competition and multiple contracting op-

portunities. Again, we assume that the items sold by the sellers or bought by buyers are equal, and each agent has only one item to buy or sell.

In the many-to-many negotiation case, each agent concurrently negotiates with many trading partners. Agent \mathbf{b}_j 's concurrent negotiation includes at most n threads $\mathfrak{S}_{\mathbf{b}_j, \mathcal{S}} = \{\mathfrak{S}_{\mathbf{b}_j, \mathbf{s}_i} | \mathbf{s}_i \in \mathcal{S}\}$, where $\mathfrak{S}_{\mathbf{b}_j, \mathbf{s}_i}$ represents the negotiation thread between \mathbf{b}_j and seller \mathbf{s}_i . We still assume that, at each time, either the buyers propose to all the sellers ($\iota(t) = \mathcal{B}$) or the sellers propose to all the buyers ($\iota(t) = \mathcal{S}$). Similarly, let $\mathcal{B}_{=t}$ be the set of buyers than t , i.e., $\mathcal{B}_{=t} = \{\mathbf{b}_j | T_{\mathbf{b}_j} = t\}$. Let \mathcal{B}_t be the set of buyers whose deadlines are not shorter deadline than t and \mathcal{B}_t^i ($\mathcal{B}_{=t}^i$) is the buyer with the i^{th} highest reserve price in \mathcal{B}_t ($\mathcal{B}_{=t}$).

We still use action *confirm* to avoid one agent's making more than one final agreement. Buyers and sellers' action space and agents' legal actions at each time are the same as that in one-to-many negotiation. The utility functions of the buyer agents are exactly those defined in the previous section. However, we need to refine the utility function of \mathbf{s}_i as it can potentially sell more items, but it has only one item to sell. We redefine \mathbf{s}_i 's utility as follows. If \mathbf{s}_i has reached more than one final agreement, it gets a utility of $-\infty$. Otherwise, it's utility is the same as that in bilateral negotiation. Therefore, \mathbf{s}_i will make at most one final agreement.

4.2.2 Agents' Equilibrium Strategies

The negotiation deadline for the negotiation between agent \mathbf{b}_j and seller \mathbf{s}_i is $T_{\mathbf{b}_j, \mathbf{s}_i} = \min(T_{\mathbf{b}_j}, T_{\mathbf{s}_i})$. The negotiation deadline for the agent \mathbf{b}_j is $T_{\mathbf{b}_j, \mathcal{S}} = \max_{\mathbf{s}_i \in \mathcal{S}} T_{\mathbf{b}_j, \mathbf{s}_i}$. Let $x_{\mathbf{b}_j, \mathbf{s}_i}^*(t)$ be \mathbf{b}_j 's optimal offer to agent \mathbf{s}_i at t if $\iota(t) = \mathcal{B}$ and $x_{\mathbf{s}_i, \mathbf{b}_j}^*(t)$ be \mathbf{s}_i 's optimal offer to agent \mathbf{b}_j at time t if $\iota(t) = \mathcal{S}$.

Lemma 15. *It is each agent's dominant strategy to propose the same price to all the trading partners at each time t .*

Proof. The proof is the same as the proof of Lemma 11. □

Then we use $x_{\mathbf{b}_j}^*(t)$ for short to represent \mathbf{b}_j 's optimal offer at t if $\iota(t) = \mathcal{B}$ and use $x_{\mathbf{s}_i}^*(t)$ to represent \mathbf{s}_i 's optimal offer at time t if $\iota(t) = \mathcal{S}$.

Lemma 16. *In equilibrium, agents of the same type should have the same equilibrium winning price (a price acceptable to agents of the different type).*

Proof. Let's prove this by contradiction. Assume two buyers have different winning prices at some time t , i.e., the lowest price acceptable to any seller. Then the seller who is willing to accept the lower winning price should change to accept the higher winning price. Therefore, the two winning prices are not in equilibrium. \square

We state the following theorem.

Theorem 17. *In the many-to-many negotiation, the sequences of optimal offers in equilibrium are: Buyer \mathbf{b}_j 's optimal offer at time $t \leq T_{\mathbf{b}_j} - 2$ is $x_{\mathbf{b}_j}^*(t) = \min(\text{RP}_{\mathbf{b}_j}, x_{\mathcal{B}}^*(t))$. Seller \mathbf{s}_i 's optimal offer at $t \leq T_{\mathbf{s}_i} - 2$ is $x_{\mathbf{s}_i}^*(t) = \max(\text{RP}_{\mathbf{s}_i}, x_{\mathcal{S}}^*(t))$.*

$x_{\mathcal{B}}^*(t)$ is given by: 1) At $t = T - 2$, $x_{\mathcal{B}}^*(t) = \text{RP}_{\mathcal{S}_{t+2}^{|\mathcal{B}_{t+2}|}}$ if $|\mathcal{B}_{t+2}| \leq |\mathcal{S}_{t+2}|$; otherwise, $x_{\mathcal{B}}^*(t) = \text{RP}_{\mathcal{B}_{t+2}^{|\mathcal{S}_{t+2}|+1}}$. 2) At $t < T - 2$, $x_{\mathcal{B}}^*(t) = \max \left\{ \text{RP}_{\mathcal{B}_{t+2}^{|\mathcal{S}_{t+2}|+1}}, \left\{ \{(x_{\mathbf{s}_i}^*(t+1))_{\leftarrow \mathbf{s}_i} | \mathbf{s}_i \in \mathcal{S}_{t+3}\} \cup \{\text{RP}_{\mathbf{s}_i} | \mathbf{s}_i \in \mathcal{S}_{t+2} - \mathcal{S}_{t+3}\}\}_{|\mathcal{S}_{t+2}|} \right\} \right.$ if $|\mathcal{S}_{t+2}| < |\mathcal{B}_{t+2}|$. *Otherwise, $x_{\mathcal{B}}^*(t) = \left\{ \{(x_{\mathbf{s}_i}^*(t+1))_{\leftarrow \mathbf{s}_i} | \mathbf{s}_i \in \mathcal{S}_{t+3}\} \cup \{\text{RP}_{\mathbf{s}_i} | \mathbf{s}_i \in \mathcal{S}_{t+2} - \mathcal{S}_{t+3}\}\}_{|\mathcal{B}_{t+2}|} \right.$ In the above equations, \mathcal{Y}_i (\mathcal{Y}^i) denotes the i^{th} smallest (largest) value in the value set \mathcal{Y} .*

$x_{\mathcal{S}}^*(t)$ is given by: 1) At $t = T - 2$, $x_{\mathcal{S}}^*(t) = \text{RP}_{\mathcal{B}_T^{|\mathcal{S}_{t+2}|}}$ if $|\mathcal{S}_T| \leq |\mathcal{B}_T|$, $x_{\mathcal{S}}^*(t) = \text{RP}_{\mathcal{S}_{t+2}^{|\mathcal{B}_{t+2}|+1}}$ if $|\mathcal{S}_{t+2}| > |\mathcal{B}_{t+2}|$. 2) At $t < T - 2$, $x_{\mathcal{S}}^*(t) = \left\{ \{(x_{\mathbf{b}_j}^*(t+1))_{\leftarrow \mathbf{b}_j} | \mathbf{b}_j \in \mathcal{B}_{t+3}\} \cup \{\text{RP}_{\mathbf{b}_j} | \mathbf{b}_j \in \mathcal{B}_{=t+2}\} \right\}^{|\mathcal{S}_{t+2}|}$ if $|\mathcal{S}_{t+2}| \leq |\mathcal{B}_{t+2}|$. *Otherwise, $x_{\mathcal{S}}^*(t) = \min \left\{ \text{RP}_{\mathcal{S}_{t+2}^{|\mathcal{B}_{t+2}|+1}}, \left\{ \{(x_{\mathbf{b}_j}^*(t+1))_{\leftarrow \mathbf{b}_j} | \mathbf{b}_j \in \mathcal{B}_{t+3}\} \cup \{\text{RP}_{\mathbf{b}_j} | \mathbf{b}_j \in \mathcal{B}_{=t+2}\} \right\}^{|\mathcal{B}_{t+2}|} \right\}$.*

Based on $x_{\mathbf{b}_j}^*(t)$ and $x_{\mathbf{s}_i}^*(t)$, we can get agents' optimal actions in the same way as that in Theorem 12 except that an agent needs to use the following rule while accepting offers or confirming accepts: a buyer \mathbf{b}_j accepts the offer x made by \mathbf{s}_i at t if: $x \leq (x_{\mathbf{b}_j}^*(t))_{\leftarrow \mathbf{b}_j}$ and x is the lowest received offer. If more than one seller has offered x and buyer \mathbf{b}_j has the q^{th} highest reserve price in \mathcal{B}_{t+2} , \mathbf{b}_j accepts the offer

made by the seller with the q^{th} lowest reserve price in sellers \mathcal{S}_{t+2} .² Similarly, if buyer \mathbf{b}_j intends to confirm an agreement with price x and multiple sellers have made the same agreement, \mathbf{b}_j will confirm the agreement made by the seller with the q^{th} lowest reserve price in sellers \mathcal{S}_{t+2} . To save space, the details of sellers' optimal actions are omitted here.

Proof. Given Lemma 15 and Lemma 16, we just need to find out the agents' equilibrium winning price at each time point. Let $x_{\mathcal{B}}^*(t)$ ($x_{\mathcal{S}}^*(t)$) be \mathcal{B} 's lowest (\mathcal{S} 's highest) offer which is acceptable to \mathcal{S} (\mathcal{B}) at time t if $\iota(t) = \mathcal{B}$ ($\iota(t) = \mathcal{S}$). It follows that $x_{\mathcal{B}}^*(t) = \max_{\mathbf{b}_j \in \mathcal{B}_{t+2}} x_{\mathbf{b}_j}^*(t)$, $x_{\mathcal{S}}^*(t) = \min_{\mathbf{s}_i \in \mathcal{S}_{t+2}} x_{\mathbf{s}_i}^*(t)$.

Following the idea of backward induction, at $T = \max_{\mathbf{b}_j \in \mathcal{B}} T_{\mathbf{b}_j, \mathcal{S}}$, the game for all agents rationally stops. The equilibrium outcome of every subgame starting from $t \geq T$ is *NoAgreement*. Therefore, at $t = T$, agents $\iota(T)$ would only confirm the best agreement proposed by agents $\iota(T-1)$. At time $t = T-1$, agents $\iota(T-1)$ will accept the best offer by agents $\iota(T-2)$ if the best offer is no worse than *NoAgreement* by accepting the best offer. At time $T-1$ and T , no agent will propose a price as it takes at least three time points to implement a final contract.

At time $t = T-2$, agents $\iota(t)$ will strive to make the best offer. There are two situations: $\iota(t) = \mathcal{B}$ or $\iota(t) = \mathcal{S}$. First consider the case $\iota(t) = \mathcal{B}$ and there are two cases: Case 1 ($|\mathcal{B}_T| \leq |\mathcal{S}_T|$): In this case, the supply is no less than demand and buyers have more bargaining power as compared with sellers. It is easy to see that each buyer's optimal price is $\text{RP}_{\mathcal{S}_T^{|\mathcal{B}_T|}}$ as, by doing so, $|\mathcal{B}_T|$ sellers will agree to sell their good and each buyer can get a good. If one buyer pays less than $\text{RP}_{\mathcal{S}_T^{|\mathcal{B}_T|}}$, the sellers will choose another buyers paying $\text{RP}_{\mathcal{S}_T^{|\mathcal{B}_T|}}$. It doesn't make sense that a rational agent will pay more than $\text{RP}_{\mathcal{S}_T^{|\mathcal{B}_T|}}$. If each buyer pays a price less than $\text{RP}_{\mathcal{S}_T^{|\mathcal{B}_T|}}$, each

²Note that in equilibrium, when a buyer \mathbf{b}_j with q^{th} highest reserve price is accepting an offer with price x , the number of sellers proposing x at $t-1$ should be no less than q . The proof is omitted as it can be easily derived from the process of calculating agents' optimal prices.

buyer will face a risk of losing an agreement as the number of sellers who are willing to accept the price is less than the number of buyers. Case 2 ($|\mathcal{B}_T| > |\mathcal{S}_T|$): In this case, the supply is less than demand and buyers need to compete with each other to get agreements. It is easy to say that each buyer's optimal price is $\text{RP}_{\mathcal{B}_T^{|\mathcal{S}_T|+1}}$. In the same way, we can get the optimal offer of buyers \mathcal{S}_T at time $T - 2$: $x_{\mathcal{S}}^*(T - 2) = \text{RP}_{\mathcal{B}_T^{|\mathcal{S}_T|}}$ if $|\mathcal{S}_T| \leq |\mathcal{B}_T|$, $x_{\mathcal{S}}^*(T - 2) = \text{RP}_{\mathcal{S}_T^{|\mathcal{B}_T|+1}}$ if $|\mathcal{S}_T| > |\mathcal{B}_T|$.

Then we move to the calculation for computing $x_{\mathcal{B}}^*(t)$ and $x_{\mathcal{S}}^*(t)$ given $x_{\mathcal{B}}^*(t + 1)$ and $x_{\mathcal{S}}^*(t + 1)$. First consider the situation that $\iota(t) = \mathcal{B}$. There are two situations depending on whether there are agents with deadline $t + 2$. If there is no agent with deadline $t + 2$, $(x_{\mathbf{s}_i}^*(t + 1))_{\leftarrow \mathbf{s}_i}$ is surely acceptable to seller \mathbf{s}_i at time t . Here we consider two cases: 1) $|\mathcal{S}_{t+3}| \geq |\mathcal{B}_{t+3}|$. It is easy to see that, the price $\min_{\mathbf{s}_i \in \mathcal{S}_{t+3}} ((x_{\mathbf{s}_i}^*(t + 1))_{\leftarrow \mathbf{s}_i})$ is surely acceptable to sellers in \mathcal{S}_{t+3} whose optimal price is $x_{\mathcal{S}}^*(t + 1)$ at time $t + 1$. However, we also need to consider the competition among buyers. Therefore, $x_{\mathcal{B}}^*(t) = \{(x_{\mathbf{s}_i}^*(t + 1))_{\leftarrow \mathbf{s}_i} | \mathbf{s}_i \in \mathcal{S}_{t+3}\}_{|\mathcal{B}_{t+3}|}$ where \mathcal{Y}_i (\mathcal{Y}^i) is the i^{th} smallest (largest) value in the value set \mathcal{Y} . 2) $|\mathcal{S}_{t+3}| < |\mathcal{B}_{t+3}|$. As $(x_{\mathbf{s}_i}^*(t + 1))_{\leftarrow \mathbf{s}_i} \leq x_{\mathbf{s}_i}^*(t + 1)$, $x_{\mathcal{B}}^*(t)$ should be no less than $\text{RP}_{\mathcal{B}_{t+3}^{|\mathcal{S}_{t+3}|+1}}$. Therefore, it follows that $x_{\mathcal{B}}^*(t) = \max \{ \text{RP}_{\mathcal{B}_{t+3}^{|\mathcal{S}_{t+3}|+1}}, \{(x_{\mathbf{s}_i}^*(t + 1))_{\leftarrow \mathbf{s}_i} | \mathbf{s}_i \in \mathcal{S}_{t+3}\}_{|\mathcal{S}_{t+3}|} \}$.

Now we move to the general case that there are some buyers or sellers with deadline $t + 2$. For a buyer with deadline $t + 2$, it is willing to propose its reserve price. For a seller with deadline $t + 2$, it is willing to accept an offer of its reserve price. Assume that there are only some sellers with deadline $t + 2$. We consider three cases: 1) $|\mathcal{S}_{t+3}| \geq |\mathcal{B}_{t+3}|$, which implies that $|\mathcal{S}_{t+2}| > |\mathcal{B}_{t+2}|$. It is easy to see that, $x_{\mathcal{B}}^*(t) = \{ \{(x_{\mathbf{s}_i}^*(t + 1))_{\leftarrow \mathbf{s}_i} | \mathbf{s}_i \in \mathcal{S}_{t+3}\} \cup \{ \text{RP}_{\mathbf{s}_i} | \mathbf{s}_i \in \mathcal{S}_{t+2} - \mathcal{S}_{t+3} \} \}_{|\mathcal{B}_{t+2}|}$. 2) $|\mathcal{S}_{t+3}| < |\mathcal{B}_{t+3}|$ and $|\mathcal{S}_{t+2}| < |\mathcal{B}_{t+2}|$. In this case, $x_{\mathcal{B}}^*(t) = \max \{ \text{RP}_{\mathcal{B}_{t+2}^{|\mathcal{S}_{t+2}|+1}}, \{ \{(x_{\mathbf{s}_i}^*(t + 1))_{\leftarrow \mathbf{s}_i} | \mathbf{s}_i \in \mathcal{S}_{t+3}\} \cup \{ \text{RP}_{\mathbf{s}_i} | \mathbf{s}_i \in \mathcal{S}_{t+2} - \mathcal{S}_{t+3} \} \}_{|\mathcal{S}_{t+2}|} \}$. 3) $|\mathcal{S}_{t+3}| < |\mathcal{B}_{t+3}|$ and $|\mathcal{S}_{t+2}| \geq |\mathcal{B}_{t+2}|$. In this case, $x_{\mathcal{B}}^*(t) = \{ \{(x_{\mathbf{s}_i}^*(t + 1))_{\leftarrow \mathbf{s}_i} | \mathbf{s}_i \in \mathcal{S}_{t+3}\} \cup \{ \text{RP}_{\mathbf{s}_i} | \mathbf{s}_i \in \mathcal{S}_{t+2} - \mathcal{S}_{t+3} \} \}_{|\mathcal{B}_{t+2}|}$.

We can easily extend the above analysis to more general cases where there are both buyers and sellers with deadline $t + 2$. We can get \mathcal{B} 's optimal price at time $t < T - 2$ as follows: 1) if $|\mathcal{S}_{t+2}| < |\mathcal{B}_{t+2}|$, $x_{\mathcal{B}}^*(t) = \max \{ \text{RP}_{\mathcal{B}}^{|\mathcal{S}_{t+2}|+1}, \{ \{ (x_{\mathbf{s}_i}^*(t+1))_{\leftarrow \mathbf{s}_i} | \mathbf{s}_i \in \mathcal{S}_{t+3} \} \cup \{ \text{RP}_{\mathbf{s}_i} | \mathbf{s}_i \in \mathcal{S}_{t+2} - \mathcal{S}_{t+3} \} \}_{|\mathcal{S}_{t+2}|} \}$; 2) otherwise, $x_{\mathcal{B}}^*(t) = \{ \{ (x_{\mathbf{s}_i}^*(t+1))_{\leftarrow \mathbf{s}_i} | \mathbf{s}_i \in \mathcal{S}_{t+3} \} \cup \{ \text{RP}_{\mathbf{s}_i} | \mathbf{s}_i \in \mathcal{S}_{t+2} - \mathcal{S}_{t+3} \} \}_{|\mathcal{B}_{t+2}|}$.

In the same way, we can get \mathcal{S} 's optimal price at time $t < T - 2$ as follows: 1) if $|\mathcal{S}_{t+2}| \leq |\mathcal{B}_{t+2}|$, $x_{\mathcal{S}}^*(t) = \{ \{ x_{\mathbf{b}_j}^*(t+1)_{\leftarrow \mathbf{b}_j} | \mathbf{b}_j \in \mathcal{B}_{t+3} \} \cup \{ \text{RP}_{\mathbf{b}_j} | \mathbf{b}_j \in \mathcal{B}_{=t+2} \} \}^{|\mathcal{S}_{t+2}|}$; 2) otherwise, $x_{\mathcal{S}}^*(t) = \min \{ \text{RP}_{\mathcal{S}}^{|\mathcal{B}_{t+2}|+1}, \{ \{ (x_{\mathbf{b}_j}^*(t+1))_{\leftarrow \mathbf{b}_j} | \mathbf{b}_j \in \mathcal{B}_{t+3} \} \cup \{ \text{RP}_{\mathbf{b}_j} | \mathbf{b}_j \in \mathcal{B}_{=t+2} \} \}^{|\mathcal{B}_{t+2}|} \}$. \square

The computational complexity of the backward induction is $\mathcal{O}((n+m)T)$ as the backward induction will go through all the time points and at each time point, each agent has at most $n+m$ possible optimal actions. It is easy to see that the bargaining agreement in the many-to-many negotiation is $(x_{\mathcal{B}}^*(0), 2)$ if $\iota(0) = \mathcal{B}$ and is $(x_{\mathcal{S}}^*(0), 2)$ if $\iota(0) = \mathcal{S}$. In addition, when the number of buyers is not equal to the number of sellers, the market competition affects the equilibrium price in the following way: if the number of buyers is less than the number of sellers, the buyers have larger bargaining power which increases with the number of sellers and decreases with the number of buyers. In contrast, if the number of buyers is larger than the number of sellers, the buyers have less bargaining power. The proposing order also affects the equilibrium price.

Figure 4.2 shows an example of backward induction construction in many-to-many negotiation. The setting in Figure 4.2 is the same as that in Figure 4.1 except that there is another buyer \mathbf{b}' with parameters $\text{RP}_{\mathbf{b}'} = 0.9$, $\delta_{\mathbf{b}'} = 0.7$, and $T_{\mathbf{b}'} = 6$. We report in the figure for any time t the optimal offer $x_{\mathcal{B}}^*(t)$ or $x_{\mathcal{S}}^*(t)$. At time $t = 4$, \mathbf{b}' can make an offer to compete with buyer \mathbf{b} . Thus we have $x_{\mathcal{B}}^*(4) = \{ (x_{\mathbf{s}_1}^*(5))_{\leftarrow \mathbf{s}_1}, (x_{\mathbf{s}_2}^*(5))_{\leftarrow \mathbf{s}_2} \}_2 = \{0.14, 0.2\}_2 = 0.2$. The process continues to the initial time point $t = 0$ where $x_{\mathcal{B}}^*(0) = 0.40992$.




Figure 4.2. Backward induction construction. At each time t the optimal offer $x_B^*(t)$ or $x_S^*(t)$ is marked.

While there is two-sided competition in the market, market mechanisms like double auction can be used for resource allocation. The double auction is one of the most common exchange institutions where both sellers and buyers submit bids which are then ranked highest to lowest to generate demand and supply profiles. Double auctions permit multiple buyers and sellers to bid to exchange a designated commodity. Some double auction mechanisms (e.g., BBDA [56]) have been applied to trading in markets. A market mechanism is *efficient* if the goods are transferred to agents that value them most.

Theorem 18. *The many-to-many negotiation is efficient.*

Proof. This result is straightforward. Assume there are sellers s_i and s_j such that $RP_{s_i} > RP_{s_j}$. It is impossible that seller s_i makes an agreement but seller s_j fails as seller s_j can make an offer lower than RP_{s_i} and thus gains a contract with positive revenue. \square

In a market consisting of two sets of agents, matching algorithms can also be used to solve agents' conflicts of resource requirements. Then we require a matching to be *stable*, i.e., it left no pair of agents on opposite sides of the market who were not matched to each other but would both prefer to be. Many-to-many negotiation allows one to avoid studying matching mechanisms since each agent is implicitly matched with all its trading partners.

4.2.3 Considerations on Settings with Uncertain Information

In this section we provide some considerations on the preliminary analysis of many-to-many bargaining with uncertainty over agents' parameters. The result discussed in Section 4.1.3 can be treated as a special case for many-to-many bargaining. With more buyers, the agreement price will increase due to the increasing competition between buyers. For the bargaining between buyers \mathcal{B} and sellers \mathcal{S} , it can be found from Theorem 17 that the agreement price depends on the reserve price of at least $\min\{|\mathcal{B}|, |\mathcal{S}|\}$ buyers and $\min\{|\mathcal{B}|, |\mathcal{S}|\}$ sellers. Although the many-to-many bargaining setting is intrinsically very complicated, the problem of finding the equilibrium outcome can be drastically simplified in some special cases.

Theorem 19. *In the following many-to-many bargaining scenarios in which $|\mathcal{B}| < |\mathcal{S}|$, the negotiation outcome only depends on the parameters of \mathcal{B} and at most $|\mathcal{B}| + 1$ sellers:*

1. *The sellers having a reserve price no higher than the $\text{RP}_{\mathcal{S}_2^{|\mathcal{B}|+1}}$ have the same deadline T' such that $\iota(T') = \mathcal{S}$.*
2. *At each time t , the seller set \mathcal{S}_{t+2} includes all the sellers with a reserve price no higher than $\text{RP}_{\mathcal{S}_2^{|\mathcal{B}_{t+2}|+1}}$.*

Proof. Case 1: At time $T' - 2$, the value of $x_{\mathcal{S}}^*(T' - 2)$ should be no higher than $\text{RP}_{\mathcal{S}_2^{|\mathcal{B}|+1}}$ and is independent of the reserve prices of sellers having a reserve price

higher than $\text{RP}_{\mathcal{S}_2^{|\mathcal{B}|+1}}$. At time $t = T' - 3$, the value of $x_{\mathcal{B}}^*(t)$ will also be no higher than $\text{RP}_{\mathcal{S}_2^{|\mathcal{B}|+1}}$. Recursively, we can find that the value of $x_{\mathcal{B}}^*(t)$ at time $t < T' - 3$ will be no higher than $\text{RP}_{\mathcal{S}_2^{|\mathcal{B}|+1}}$ and is independent of the reserve prices of sellers having a reserve price higher than $\text{RP}_{\mathcal{S}_2^{|\mathcal{B}|+1}}$.

Case 2: We can prove the result in the same way as in the proof of *Case 1*. \square

Thus, the negotiation outcome only depends on a small number of parameters in some special cases. The complexity of solving complete information bargaining and incomplete information bargaining can be reduced.

4.3 Uncertainty about Reserve Prices

In this section we analyze agents' rational strategies in concurrent negotiation with incomplete information. More specifically, we focus on the situation that one buyer \mathbf{b} is negotiating with a number of sellers \mathcal{S} and there is uncertainty about the buyer's reserve price. We extend our algorithm for bilateral bargaining to handle concurrent negotiation.

4.3.1 Introducing Uncertainty

We assume the one-sided uncertainty regarding the type of the buyer \mathbf{b} (the case of having uncertainty with the type of a seller $\mathbf{s} \in \mathcal{S}$ can be analyzed analogously). The buyer \mathbf{b} can be of finitely many types $\{\mathbf{b}_1, \dots, \mathbf{b}_n\}$ in which buyer \mathbf{b}_i has a reserve price RP_i . The initial belief of \mathbf{s} on \mathbf{b} is $\mu(0) = \langle \Delta_{\mathbf{b}}^0, P_{\mathbf{b}}^0 \rangle$ where $\Delta_{\mathbf{b}}^0 = \{\mathbf{b}_1, \dots, \mathbf{b}_n\}$ and $P_{\mathbf{b}}^0 = \{\omega_{\mathbf{b}_1}^0, \dots, \omega_{\mathbf{b}_n}^0\}$ such that $\sum_i \omega_{\mathbf{b}_i}^0 = 1$. $\omega_{\mathbf{b}_i}^0$ is the *priori* probability that \mathbf{b} is of type \mathbf{b}_i . During bargaining, seller \mathbf{s} 's belief will evolve using the Bayes rule. The belief of \mathbf{s} on the type of \mathbf{b} at time t is $\mu(t)$. It's easy to see that in incomplete information bargaining, it's still a weekly dominant strategy for the buyer \mathbf{b} to make the same offer to all the sellers. Therefore, different sellers' beliefs about the type of buyer \mathbf{b} will always be the same at any time t . The belief of \mathbf{s} on the type of \mathbf{b} at

time t is $\mu(t)$. The probability assigned by \mathbf{s} to $\mathbf{b} = \mathbf{b}_i$ at time t is denoted $\omega_{\mathbf{b}_i}^t$. Given an assessment $a = \langle \mu, \sigma \rangle$, there are multiple possible bargaining outcomes: outcome $o_{\mathbf{b}_i}$ if $\mathbf{b} = \mathbf{b}_i$. We denote bargaining outcome as $o = \langle o_{\mathbf{b}_1}, \dots, o_{\mathbf{b}_n} \rangle$.

With pure strategies, buyer types' possible behaviors regarding whether they behave in the same way on the equilibrium path at each decision making node are finite. We use the term “choice rule” to characterize buyer types' strategies regarding whether they behave in the same way at a specific decision making point. Easily, at a decision making node \mathbf{b}_i and \mathbf{b}_j can make the same offer (in this case, choice rules are said *pooling*) or can make different offers (in this case, choice rules are said *separating*). On the basis of this consideration, we can make some assumptions over the belief system without losing generality. On the equilibrium path $\mu(t) = \langle \Delta_{\mathbf{b}}^t, P_{\mathbf{b}}^t \rangle$ of \mathbf{s} on \mathbf{b} at any time t is one the following. After a time point t where buyer types' choice rule is pooling, $\mu(t+1) = \mu(t)$, i.e., $\Delta_{\mathbf{b}}^{t+1} = \Delta_{\mathbf{b}}^t$ and $P_{\mathbf{b}}^{t+1} = P_{\mathbf{b}}^t$. As is customary in economic studies [112], we consider only stationary systems of beliefs, i.e., if a seller \mathbf{s} believes a \mathbf{b} 's type with zero probability at time point t , then it will continue to believe such a type with zero probability at any time point $t' > t$. We need also specify the belief system off the equilibrium path, i.e., when an agent makes an action that is not optimal. We use the *optimistic conjectures* [112]. That is, when \mathbf{b} acts off the equilibrium strategy, agent \mathbf{s} will believe that agent \mathbf{b} is of its “weakest” type, i.e., the type against which each seller would gain the most. In our case, the weakest type is the buyer type with the highest reserve price (we prove it in the following section). We can therefore specify $\mu(t)$ by specifying $\Delta_{\mathbf{b}}^t$. That is, if $\mu(t-1) = \Delta_{\mathbf{b}}^{t-1}$ and \mathbf{b} acts off the equilibrium strategy at time $t-1$, it follows that $\Delta_{\mathbf{b}}^t = \mathbf{b}_h(\Delta_{\mathbf{b}}^{t-1})$ where $\mathbf{b}_h(\Delta_{\mathbf{b}}^{t-1})$ is the buyer type with the highest reserve price in buyer types $\Delta_{\mathbf{b}}^{t-1}$.

4.3.2 Off the Equilibrium Path Optimal Strategies

Before analyzing equilibrium strategies when the buyer can be of many types, we provide the optimal strategies in the situations \mathbf{s} believes the buyer of one single type. There are two cases: 1) Seller \mathbf{s} has the right belief about the type of the buyer \mathbf{b} . In this case, agents' equilibrium strategies are the equilibrium strategies of the corresponding complete information bargaining discussed in Section 4.1. Let $x_{\mathbf{b}_i}^c(t)$ be agents' optimal offer at time t when \mathbf{b} is of type \mathbf{b}_i in this case. That is, if $\iota(t) = \mathbf{b}$, $x_{\mathbf{b}_i}^c(t)$ is \mathbf{b} 's optimal offer $x_{\mathbf{b}}^*(t)$ at time t in complete information bargaining when it is of type \mathbf{b}_i . 2) Seller \mathbf{s} has the wrong belief about the type of the buyer \mathbf{b} , i.e., \mathbf{b}_i is believed to be \mathbf{b}_j and \mathbf{b}_j is believed to be \mathbf{b}_i .

Lemma 20. $x_{\mathbf{b}_i}^c(t) \geq x_{\mathbf{b}_j}^c(t)$ if $\text{RP}_i > \text{RP}_j$.

Proof. We can prove the results from the proof of Theorem 12:

Case 1 ($\iota(T) = \mathbf{s}$). It follows that $x_{\mathbf{b}_i}^c(T-2) = x_{\mathbf{b}_j}^c(T-2) = \text{RP}_{S_T^2}$. Then we have $x_{\mathbf{b}_i}^c(T-3) = \min\{(x_{\mathbf{b}_i}^c(T-2))_{\leftarrow S_{T-1}^1}, \text{RP}_{S_{T-1}^2}\} = \min\{(x_{\mathbf{b}_j}^c(T-2))_{\leftarrow S_{T-1}^1}, \text{RP}_{S_{T-1}^2}\} = x_{\mathbf{b}_j}^c(T-3)$. At time $t = T-4$, we have $x_{\mathbf{b}_i}^c(t) = \min\{\text{RP}_{S_{t+2}^2}, (x_{\mathbf{b}_i}^c(t+1))_{\leftarrow \mathbf{b}}\} = \min\{\text{RP}_{S_{t+2}^2}, \text{RP}_i(1 - \delta_{\mathbf{b}}) + \delta_{\mathbf{b}}x_{\mathbf{b}_i}^c(t+1)\} \geq \min\{\text{RP}_{S_{t+2}^2}, \text{RP}_i(1 - \delta_{\mathbf{b}}) + \delta_{\mathbf{b}}x_{\mathbf{b}_j}^c(t+1)\} = x_{\mathbf{b}_j}^c(t)$. Recursively, we have $x_{\mathbf{b}_i}^c(t) \geq x_{\mathbf{b}_j}^c(t)$ for $t < T-4$.

Case 2 ($\iota(T) = \mathbf{b}$). It follows that $x_{\mathbf{b}_i}^c(T-2) = \text{RP}_{S_T^1} = x_{\mathbf{b}_j}^c(T-2)$. Then at time $T-3$, we have $x_{\mathbf{b}_i}^c(T-3) = \min\{\text{RP}_{S_{T-1}^2}, (x_{\mathbf{b}_i}^c(T-2))_{\leftarrow \mathbf{b}}\} = \min\{\text{RP}_{S_{T-1}^2}, \text{RP}_i(1 - \delta_{\mathbf{b}}) + \delta_{\mathbf{b}}x_{\mathbf{b}_i}^c(T-2)\} \geq \min\{\text{RP}_{S_{T-1}^2}, \text{RP}_i(1 - \delta_{\mathbf{b}}) + \delta_{\mathbf{b}}x_{\mathbf{b}_j}^c(T-2)\} = x_{\mathbf{b}_j}^c(T-3)$. Recursively, we have $x_{\mathbf{b}_i}^c(t) \geq x_{\mathbf{b}_j}^c(t)$ for $t < T-3$. \square

We can see that \mathbf{b}_i is weaker than \mathbf{b}_j in terms of its offering price at each time point in complete information bargaining. Furthermore, we can get $\text{RP}_i - x_{\mathbf{b}_i}^c(t) \geq \text{RP}_j - x_{\mathbf{b}_j}^c(t)$ following the same procedure in the proof of Lemma 20. $\text{RP}_i - x_{\mathbf{b}_i}^c(0)$ is the gain (utility) of \mathbf{b}_i in complete information bargaining and $\text{RP}_j - x_{\mathbf{b}_j}^c(0)$ is the gain (utility) of \mathbf{b}_j in complete information bargaining.

Lemma 21. *If $RP_i > RP_j$, it follows that $x_{\mathbf{b}_i}^c(t) \leq (x_{\mathbf{b}_i}^c(t+1))_{\leftarrow \mathbf{b}_i}$ and $x_{\mathbf{b}_j}^c(t) \leq (x_{\mathbf{b}_j}^c(t+1))_{\leftarrow \mathbf{b}_j}$.*

Proof. We can get this result by following the same procedure in the proof of Lemma 20. This result indicates that the buyer will accept sellers' lowest equilibrium price in complete information bargaining, i.e., agents will reach a final agreement at time $t - 2$ in complete information bargaining. \square

Agents' optimal strategies when any seller \mathbf{s} has the wrong belief about the type of the buyer \mathbf{b} are shown in the following theorem:

Theorem 22. *If seller \mathbf{s} has the wrong belief about the type of \mathbf{b} , the optimal strategies of any seller \mathbf{s} are those in complete information bargaining. The optimal strategies $\sigma_{\mathbf{b}_i}^*(t)|\{\mathbf{b}_j\}$ of buyer \mathbf{b}_i when it's believed to be \mathbf{b}_j are:*

$$\sigma_{\mathbf{b}_i}^*(t)|\{\mathbf{b}_j\} = \begin{cases} \text{accept } y & \text{if } y \leq (x_{\mathbf{b}_j}^c(t))_{\leftarrow \mathbf{b}_i} \\ \text{offer } x_{\mathbf{b}_j}^c(t) & \text{otherwise} \end{cases}$$

The optimal strategies $\sigma_{\mathbf{b}_j}^(t)|\{\mathbf{b}_i\}$ of the buyer \mathbf{b}_j when it's believed to be \mathbf{b}_i are:*

$$\sigma_{\mathbf{b}_j}^*(t)|\{\mathbf{b}_i\} = \begin{cases} \text{accept } y & \text{if } y \leq \min\{(x_{\mathbf{b}_i}^c(t))_{\leftarrow \mathbf{b}_j}, RP_j\} \\ \text{offer } \min\{x_{\mathbf{b}_i}^c(t), RP_j\} & \text{otherwise} \end{cases}$$

Proof. Case 1 (\mathbf{b}_i is believed to be \mathbf{b}_j). If sellers' lowest offer at time $t - 2$ is $x_{\mathbf{b}_j}^c(t - 1)$, buyer \mathbf{b}_i 's optimal strategy is to accept it as the minimum price that the seller would accept at time $t + 1$, i.e., $x_{\mathbf{b}_j}^c(t)$, gives \mathbf{b}_i a utility lesser than $x_{\mathbf{b}_j}^c(t - 1)$ since $(x_{\mathbf{b}_j}^c(t))_{\leftarrow \mathbf{b}_i} > (x_{\mathbf{b}_j}^c(t))_{\leftarrow \mathbf{b}_j} \geq x_{\mathbf{b}_j}^c(t - 1)$. If the seller acts off the equilibrium path and offers a price y lower than $x_{\mathbf{b}_j}^c(t - 1)$, the optimal strategy of \mathbf{b}_i is obviously to accept y . If the seller offers a price y higher than $x_{\mathbf{b}_j}^c(t - 1)$, the optimal strategy of \mathbf{b}_i is to accept y only if $y \leq (x_{\mathbf{b}_j}^c(t))_{\leftarrow \mathbf{b}_i}$, otherwise \mathbf{b}_i 's optimal strategy is to reject y and to offer $x_{\mathbf{b}_j}^c(t)$. Note that $x_{\mathbf{b}_i}^c(t) \leq RP_i$ and $x_{\mathbf{b}_j}^c(t) \leq RP_i$.

Case 2 (\mathbf{b}_j is believed to be \mathbf{b}_i). This case is more complicated as sellers' lowest offer $x_{\mathbf{b}_i}^c(t - 1)$ at time t on its equilibrium path may be not acceptable to \mathbf{b}_j as when

\mathbf{b}_j offers $x_{\mathbf{b}_i}^c(t)$ at time t , it follows that $(x_{\mathbf{b}_i}^c(t))_{\leftarrow \mathbf{b}_j} < (x_{\mathbf{b}_i}^c(t))_{\leftarrow \mathbf{b}_i}$ and $(x_{\mathbf{b}_i}^c(t))_{\leftarrow \mathbf{b}_i} \geq x_{\mathbf{b}_i}^c(t-1)$ (Lemma 21). In addition, \mathbf{b}_j may not offer $x_{\mathbf{b}_i}^c(t)$ if $x_{\mathbf{b}_i}^c(t)$ is higher than RP_j . Therefore, \mathbf{b}_j 's optimal offer at time t is $\min\{x_{\mathbf{b}_i}^c(t), \text{RP}_j\}$. Thus, \mathbf{b}_j will accept an offer y at time t such that $y \leq \min\{(x_{\mathbf{b}_i}^c(t))_{\leftarrow \mathbf{b}_j}, \text{RP}_j\}$. \square

4.3.3 Our Approach

While having multiple buyers increases the complexity of computing sequential equilibria, we can extend our approach for bilateral bargaining with uncertainty in Chapter 3 to handle one-to-many bargaining with uncertainty. When it is the buyer's turn to make an offer, we consider different choice rules. Note that the number of choice rules does not depend on the number of sellers since when buyer types are making an acceptable offer, they only need to consider the offer that is acceptable to the seller with the lowest reserve price due to the market competition between different sellers. When it is the seller's turn to make an offer, we consider different reject update rules. Due to competition between different sellers, we only need to consider the reject update rule of the seller with the lowest reserve price. However, when we compute the optimal offer for the seller or a buyer type, market competition should be taken into account. In what follows we briefly discuss how to compute agents' equilibrium offers while using the algorithm presented in Chapter 3.

4.3.4 The Buyer's Equilibrium Offer

Now we consider the buyer's equilibrium offer at a continuation game $\Gamma(t, \Delta_{\mathbf{b}})$ such that $\iota(t) = \mathbf{b}$. If $t = T - 1$, it is the buyer agent's dominant strategy to accept any offer which is not worse than its reserve price. At time $t = T - 2$, different buyer types' optimal offer is $\text{RP}_{\mathcal{S}_T^1}$ since seller \mathcal{S}_T^1 will accept the offer at time $T - 1$. If $|\Delta_{\mathbf{b}}| = 1$, agents' equilibrium strategies are the equilibrium strategies of the corresponding complete information bargaining. When $|\Delta_{\mathbf{b}}| > 1$ at time $t < T - 2$, buyer types

have multiple choice rules and we need to consider the equilibrium strategies for each choice rule.

4.3.4.1 Pooling Choice Rule

When \mathbf{b} employs a pooling choice rule at a continuation game $\Gamma(t, \Delta_{\mathbf{b}})$, seller \mathbf{s} will not change its belief after observing the buyer's equilibrium offer since all buyer types will behave in the same way. Thus, we need to consider all sequential equilibria $\mathbf{SE}(\Delta_{\mathbf{b}}, t + 1)$ of the continuation game with belief $\Delta_{\mathbf{b}}$ at time $t + 1$. If $\mathbf{SE}(\Delta_{\mathbf{b}}, t + 1) = \emptyset$, there is no sequential equilibrium for this choice rule. Otherwise, for each sequential equilibrium $\wp \in \mathbf{SE}(\Delta_{\mathbf{b}}, t + 1)$, we compute buyer types' optimal offer and check the satisfaction of equilibrium existence conditions.

First we consider the accepting pooling choice rule. Let $e_{\mathcal{S}_{t+3}^1}^{t+1} | \Delta_{\mathbf{b}}$ be \mathcal{S}_{t+3}^1 's equivalent offer at time $t + 1$ given the belief $\Delta_{\mathbf{b}}$ in the sequential equilibrium \wp . At time $t + 1$, the equilibrium strategy of \mathcal{S}_{t+3}^1 is that \mathcal{S}_{t+3}^1 will accept any offer y if $y \geq (e_{\mathcal{S}_{t+3}^1}^{t+1} | \Delta_{\mathbf{b}})_{\leftarrow \mathcal{S}_{t+3}^1}$. Note that if $T_{\mathcal{S}_{t+2}^1} = t + 2$, seller \mathcal{S}_{t+2}^1 is willing to accept any offer which is no worse than its reserve price. Therefore, the equilibrium offer of buyer $\mathbf{b}_i \in \Delta_{\mathbf{b}}$ at time t is

$$x_{\mathbf{b}_i}^*(t) | \Delta_{\mathbf{b}} = \begin{cases} \text{RP}_{\mathcal{S}_{t+2}^1} & T_{\mathcal{S}_{t+2}^1} = t + 2 \\ \min\{(e_{\mathcal{S}_{t+3}^1}^{t+1} | \Delta_{\mathbf{b}})_{\leftarrow \mathcal{S}_{t+3}^1}, \text{RP}_{\mathcal{S}_{t+2}^2}\} & T_{\mathcal{S}_{t+2}^1} \neq t + 2 \text{ and } |\mathcal{S}_{t+2}| > 1 \\ (e_{\mathcal{S}_{t+3}^1}^{t+1} | \Delta_{\mathbf{b}})_{\leftarrow \mathcal{S}_{t+3}^1} & T_{\mathcal{S}_{t+2}^1} \neq t + 2 \text{ and } |\mathcal{S}_{t+2}| = 1 \end{cases}$$

The corresponding offer update rule is the following: $\mu(t + 1) = \Delta_{\mathbf{b}}$ if $\sigma_{\mathbf{b}}(t) = \text{offer } x_{\mathbf{b}_i}^*(t) | \Delta_{\mathbf{b}}$; $\mu(t + 1) = \{\mathbf{b}_h(\Delta_{\mathbf{b}})\}$, otherwise. If buyer $\mathbf{b}_i \in \Delta_{\mathbf{b}}$ deviates from offering $x_{\mathbf{b}_i}^*(t) | \Delta_{\mathbf{b}}$, it will be believed to be of type $\mathbf{b}_h(\Delta_{\mathbf{b}})$. Following Theorem 22, when a buyer \mathbf{b}_i is believed to be of type $\mathbf{b}_h(\Delta_{\mathbf{b}})$ which has a reserve price no less than RP_i , \mathbf{b}_i 's optimal offer at time t is $x_{\mathbf{b}_i}^*(t) | \{\mathbf{b}_h(\Delta_{\mathbf{b}})\}$. Thus, the condition of

equilibrium existence needed to be checked is $x_{\mathbf{b}_i}^*(t)|\Delta_{\mathbf{b}} \leq x_{\mathbf{b}_i}^*(t)|\{\mathbf{b}_h(\Delta_{\mathbf{b}})\}$ for all $\mathbf{b}_i \in \Delta_{\mathbf{b}}$. If the equilibrium existence conditions are satisfied, there is a sequential equilibrium with buyer types' offer $x_{\mathbf{b}_i}^*(t)|\Delta_{\mathbf{b}}$ and \wp as the sequential equilibrium for the continuation game from time $t + 1$. Buyer \mathbf{b}_i 's equilibrium bargaining outcome in this equilibrium is $EBO(\mathbf{b}_i, x_{\mathbf{b}_i}^*(t)|\Delta_{\mathbf{b}}, t) = (x_{\mathbf{b}_i}^*(t)|\Delta_{\mathbf{b}}, t + 1)$ since $x_{\mathbf{b}_i}^*(t)|\Delta_{\mathbf{b}}$ is acceptable to the seller. Thus buyer \mathbf{b}_i 's equivalent offer is $e_{\mathbf{b}_i}^t|\Delta_{\mathbf{b}} = x_{\mathbf{b}_i}^*(t)|\Delta_{\mathbf{b}}$.

Next we consider the rejecting pooling choice rule where all buyer types $\Delta_{\mathbf{b}}$ will make an offer (i.e., -1) that will be rejected by the seller. Buyer \mathbf{b}_i 's equilibrium bargaining outcome is the bargaining outcome in the sequential equilibrium \wp , i.e., $EBO(\mathbf{b}_i, -1, t) = EBO(\mathbf{b}_i, \wp)$. Thus buyer \mathbf{b}_i 's equivalent offer is $e_{\mathbf{b}_i}^t|\Delta_{\mathbf{b}} = \rho(\mathbf{b}_i, t, EBO(\mathbf{b}_i, \wp))$ where function $\rho(\mathbf{b}_i, t, EBO(\mathbf{b}_i, \wp))$ which satisfies $U_{\mathbf{b}_i}(\rho(\mathbf{b}_i, t, EBO(\mathbf{b}_i, \wp)), t + 1) = U_{\mathbf{b}_i}(EBO(\mathbf{b}_i, \wp))$. The corresponding offer update rule is the following: $\mu(t + 1) = \Delta_{\mathbf{b}}$ if $\sigma_{\mathbf{b}}(t) = \text{offer} - 1$; $\mu(t + 1) = \{\mathbf{b}_h(\Delta_{\mathbf{b}})\}$, otherwise. If buyer \mathbf{b} deviates from offering -1 at time t , it will be treated as buyer type $\mathbf{b}_h(\Delta_{\mathbf{b}})$ and the equilibrium existence condition is $e_{\mathbf{b}_i}^t|\Delta_{\mathbf{b}} \leq x_{\mathbf{b}_i}^*(t)|\{\mathbf{b}_h(\Delta_{\mathbf{b}})\}$ for all $\mathbf{b}_i \in \Delta_{\mathbf{b}}$.

4.3.4.2 Separating Choice Rule

Now we consider agents' equilibrium strategies at a continuation game $\Gamma(t, \Delta_{\mathbf{b}})$ when buyer \mathbf{b} employs the separating choice rule where buyer types $\Delta_{\mathbf{b}}^a$ make an acceptable offer while buyer types $\Delta_{\mathbf{b}}^r$ make a rejectable offer -1 . For this choice rule, the reasonable beliefs of its continuation game are $\Delta_{\mathbf{b}}^a$ and $\Delta_{\mathbf{b}}^r$. If one of the continuation games has no sequential equilibrium, there is no sequential equilibrium for this choice rule. In what follows we show how to compute agents' equilibrium strategies at time t given a sequential equilibrium $\wp^a \in \mathbf{SE}(\Delta_{\mathbf{b}}^a, t + 1)$ and a sequential equilibrium $\wp^r \in \mathbf{SE}(\Delta_{\mathbf{b}}^r, t + 1)$.

Let $e_{\mathcal{S}_{t+3}^1}^{t+1}|\Delta_{\mathbf{b}}^a$ be \mathbf{s} 's equivalent offer at time $t+1$ in the equilibrium \wp^a . Let $e_{\mathcal{S}_{t+3}^1}^{t+1}|\Delta_{\mathbf{b}}^r$ ($x^{t+1}|\Delta_{\mathbf{b}}^r$, respectively) be \mathcal{S}_{t+3}^1 's equivalent offer (equilibrium offer, respectively) at time $t+1$ in the equilibrium \wp^r . By convention, the equilibrium offer of buyer type $\mathbf{b}_j \in \Delta_{\mathbf{b}}^r$ at time t is -1 . Buyer \mathbf{b}_j 's equilibrium bargaining outcome is the bargaining outcome $EBO(\mathbf{b}_j, \wp^r)$ in the sequential equilibrium \wp^r . Thus buyer $\mathbf{b}_j \in \Delta_{\mathbf{b}}^r$'s equivalent offer is $e_{\mathbf{b}_j}^t|\Delta_{\mathbf{b}} = \rho(\mathbf{b}_j, t, EBO(\mathbf{b}_j, \wp^r))$. Similar to the pooling acceptance choice rule, the optimal offer of buyer types $\Delta_{\mathbf{b}}^a$ at time t is as follows

$$x_{\mathbf{b}_i}^*(t)|\Delta_{\mathbf{b}} = \begin{cases} \text{RP}_{\mathcal{S}_{t+2}^1} & T_{\mathcal{S}_{t+2}^1} = t+2 \\ \min\{(e_{\mathcal{S}_{t+3}^1}^{t+1}|\Delta_{\mathbf{b}}^a)_{\leftarrow \mathcal{S}_{t+3}^1}, \text{RP}_{\mathcal{S}_{t+2}^2}\} & T_{\mathcal{S}_{t+2}^1} \neq t+2 \text{ and } |\mathcal{S}_{t+2}| > 1 \\ (e_{\mathcal{S}_{t+3}^1}^{t+1}|\Delta_{\mathbf{b}}^a)_{\leftarrow \mathcal{S}_{t+3}^1} & T_{\mathcal{S}_{t+2}^1} \neq t+2 \text{ and } |\mathcal{S}_{t+2}| = 1 \end{cases}$$

Accordingly, buyer $\mathbf{b}_i \in \Delta_{\mathbf{b}}^a$'s equivalent offer is $e_{\mathbf{b}_i}^t|\Delta_{\mathbf{b}} = x_{\mathbf{b}_i}^*(t)|\Delta_{\mathbf{b}}$ since its equilibrium bargaining outcome is $(x_{\mathbf{b}_i}^*(t)|\Delta_{\mathbf{b}}, t+1)$. Seller \mathbf{s} will update its belief to $\Delta_{\mathbf{b}}^a$ when it receives an offer $x_{\mathbf{b}_i}^*(t)|\Delta_{\mathbf{b}}$. If it receives an offer -1 , it will update its belief to $\Delta_{\mathbf{b}}^r$. Otherwise, it will update its belief to $\mathbf{b}_h(\Delta_{\mathbf{b}})$. The existence of such an equilibrium depends on the following conditions:

- Any buyer type $\mathbf{b}_i \in \Delta_{\mathbf{b}}^a$ has no incentive to behave as any buyer type $\mathbf{b}_j \in \Delta_{\mathbf{b}}^r$. If \mathbf{b}_i pretends to be \mathbf{b}_j , it will offer -1 at time t and its equilibrium bargaining outcome will be $EBO(\mathbf{b}_j, -1, t) = EBO(\mathbf{b}_j, \wp^r)$. Therefore, this condition requires that $U_{\mathbf{b}_i}(EBO(\mathbf{b}_i, x_{\mathbf{b}_i}^*(t)|\Delta_{\mathbf{b}}, t)) \geq U_{\mathbf{b}_i}(EBO(\mathbf{b}_j, \wp^r))$ or equivalently, $x_{\mathbf{b}_i}^*(t)|\Delta_{\mathbf{b}} \leq \rho(\mathbf{b}_i, t, EBO(\mathbf{b}_j, \wp^r))$.
- Any buyer type $\mathbf{b}_j \in \Delta_{\mathbf{b}}^r$ must have no incentive to behave as $\mathbf{b}_i \in \Delta_{\mathbf{b}}^a$. If \mathbf{b}_j behaves as \mathbf{b}_i , it will offer $x_{\mathbf{b}_i}^*(t)|\Delta_{\mathbf{b}}$ at time t and the offer will be accepted. \mathbf{b}_j will not choose to behave as \mathbf{b}_i if $U_{\mathbf{b}_j}(EBO(\mathbf{b}_j, \wp^r)) \geq U_{\mathbf{b}_j}(EBO(\mathbf{b}_j, x_{\mathbf{b}_i}^*(t)|\Delta_{\mathbf{b}}, t))$ or equivalently, $\rho(\mathbf{b}_j, t, EBO(\mathbf{b}_j, \wp^r)) \leq x_{\mathbf{b}_i}^*(t)|\Delta_{\mathbf{b}}$.

- No buyer type has an incentive to offer a price different from the above two equilibrium offers. If a buyer type $\mathbf{b}_i \in \Delta_{\mathbf{b}}$ offers a price different from $x_{\mathbf{b}_i}^*(t)|\Delta_{\mathbf{b}}$ and -1 , it will be treated as buyer type $\mathbf{b}_h(\Delta_{\mathbf{b}})$ and its optimal offer at time t is then $x_{\mathbf{b}_i}^*(t)|\{\mathbf{b}_h(\Delta_{\mathbf{b}})\}$. Buyer type \mathbf{b}_i will not choose to act off the equilibrium path if $e_{\mathbf{b}_i}^t|\Delta_{\mathbf{b}} \leq x_{\mathbf{b}_i}^*(t)|\{\mathbf{b}_h(\Delta_{\mathbf{b}})\}$.

If all the three conditions are satisfied, buyer types' optimal offers, the belief update rule, and the sequential equilibria \wp^a and \wp^r for the continuation game starting from time $t + 1$ consists of a sequential equilibrium for the continuation game $\Gamma(t, \Delta_{\mathbf{b}})$.

4.3.5 The Seller's Equilibrium Offer

Now we show how to compute the seller's equilibrium offer at a continuation game $\Gamma(t, \Delta_{\mathbf{b}})$ such that $\iota(t) = \mathcal{S}$. If $t = T - 1$, it is the seller's dominant strategy to accept any offer which is not worse than its reserve price. At time $t = T - 2$, there are two cases. If $|\mathcal{S}_{t+2}| = 1$, seller \mathcal{S}_{t+2}^1 has multiple choices, each for one buyer type in $\Delta_{\mathbf{b}}$. The optimal offer of seller \mathcal{S}_{t+2}^1 for buyer type $\mathbf{b}_i \in \Delta_{\mathbf{b}}$ is RP_i , which gives seller \mathcal{S}_{t+2}^1 an expected utility $EU_{\mathcal{S}_{t+2}^1}(\text{RP}_i, t + 2) = \sum_{\mathbf{b}_j \in \Delta_{\mathbf{b}}, \text{RP}_j \geq \text{RP}_i} \omega_{\mathbf{b}_j}(\Delta_{\mathbf{b}}) U_{\mathcal{S}_{t+2}^1}(\text{RP}_i, t + 2)$ since RP_i is only acceptable to a buyer type with a reserve price no less than RP_i . The optimal offer of \mathcal{S}_{t+2}^1 at time $t = T - 2$ is $y = \arg \max_{y \in \{\text{RP}_i | \mathbf{b}_i \in \Delta_{\mathbf{b}}\}} EU_{\mathcal{S}_{t+2}^1}(y, t)$ and its equivalent offer is $e_{\mathcal{S}_{t+2}^1}^t|\Delta_{\mathbf{b}}$ such that $U_{\mathcal{S}_{t+2}^1}(e_{\mathcal{S}_{t+2}^1}^t|\Delta_{\mathbf{b}}, t + 2) = EU_{\mathcal{S}_{t+2}^1}(y, t)$. If $|\mathcal{S}_{t+2}| > 1$, seller \mathcal{S}_{t+2}^1 's optimal offer at time t is $\text{RP}_{\mathcal{S}_{t+2}^2}$ due to the competition between sellers. Thus, the equivalent price of the optimal offer of agent \mathcal{S}_{t+2}^1 in this case is $\text{RP}_{\mathcal{S}_{t+2}^2}$. If $|\Delta_{\mathbf{b}}| = 1$, agents' equilibrium strategies are the equilibrium strategies of the corresponding complete information bargaining.

Now we show how to compute the seller \mathcal{S}_{t+2}^1 's equilibrium offer given a belief $\Delta_{\mathbf{b}}$ ($|\Delta_{\mathbf{b}}| > 1$) at time $t < T - 2$. We consider a reject update rule in which buyer types $\Delta'_{\mathbf{b}}$ will reject the seller \mathcal{S}_{t+2}^1 's offer and buyer types $\Delta_{\mathbf{b}} - \Delta'_{\mathbf{b}}$ will accept the seller \mathcal{S}_{t+2}^1 's offer such that such that $\text{RP}_i > \text{RP}_j$ for any $\mathbf{b}_i \in \Delta_{\mathbf{b}} - \Delta'_{\mathbf{b}}$ and $\mathbf{b}_j \in \Delta'_{\mathbf{b}}$. We

first compute all the sequential equilibria $\mathbf{SE}(\Delta'_{\mathbf{b}}, t + 1)$ for the continuation game with belief $\Delta'_{\mathbf{b}}$ starting from time $t + 1$. If there is no sequential equilibrium for the continuation game $\Gamma(t+1, \Delta'_{\mathbf{b}})$, there is no sequential equilibrium for this reject update rule. Otherwise, for each sequential equilibrium $\varphi \in \mathbf{SE}(\Delta'_{\mathbf{b}}, t + 1)$, we check whether there exists a price x such that the price, the reject update rule, and the sequential equilibrium φ constitute a sequential equilibrium for the continuation game $\Gamma(t, \Delta_{\mathbf{b}})$. Such a price x exists if and only if the following three conditions are satisfied:

1. \mathbf{b}_i is willing to accept the offer x and does not want to behave as \mathbf{b}_j . That is, for any $\mathbf{b}_i \in \Delta_{\mathbf{b}} - \Delta'_{\mathbf{b}}$ and $\mathbf{b}_j \in \Delta'_{\mathbf{b}}$, $U_{\mathbf{b}_i}(x, t + 2) \geq U_{\mathbf{b}_i}(EBO(\mathbf{b}_j, \varphi))$ where $EBO(\mathbf{b}_j, \varphi)$ is the \mathbf{b}_i 's equilibrium bargaining outcome when it behaves as \mathbf{b}_j . This condition can be reformulated as $x \leq \min_{\mathbf{b}_i \in \Delta_{\mathbf{b}} - \Delta'_{\mathbf{b}}, \mathbf{b}_j \in \Delta'_{\mathbf{b}}} \rho(\mathbf{b}_i, t, EBO(\mathbf{b}_j, \varphi))$, which provides an upper bound for seller's offering price x .
2. \mathbf{b}_j will reject the offer x . That is, each buyer type $\mathbf{b}_j \in \Delta'_{\mathbf{b}}$ has no incentive to behave as \mathbf{b}_i , i.e., $U_{\mathbf{b}_j}(x, t + 2) < U_{\mathbf{b}_j}(EBO(\mathbf{b}_j, \varphi))$. This condition can be rewritten as $x > \max_{\mathbf{b}_j \in \Delta'_{\mathbf{b}}} \rho(\mathbf{b}_j, t, EBO(\mathbf{b}_j, \varphi))$, which provides a lower bound for the offering price x .
3. Seller \mathbf{s} has no incentive to choose a price other than x given the reject update rule and the sequential equilibrium φ of the continuation game $\Gamma(t + 1, \Delta'_{\mathbf{b}})$;

The third equilibrium existence condition requires that the price x is seller \mathcal{S}_{t+2}^1 's optimal offer given the reject update rule and the sequential equilibrium φ for the continuation game. Any buyer type can either accept the seller \mathcal{S}_{t+2}^1 's offer x or reject it and receive a bargaining outcome in the sequential equilibrium φ for the continuation game. Formally, buyer type $\mathbf{b}_j \in \Delta'_{\mathbf{b}}$ will accept a price x if and only if $x \leq \rho(\mathbf{b}_j, t, EBO(\mathbf{b}_j, \varphi))$. Buyer type $\mathbf{b}_i \in \Delta_{\mathbf{b}} - \Delta'_{\mathbf{b}}$ will accept a price x if and only if $x \leq \min_{\mathbf{b}_j \in \Delta'_{\mathbf{b}}} \rho(\mathbf{b}_i, t, EBO(\mathbf{b}_j, \varphi))$. We can define the acceptance price $\phi(\mathbf{b}_i, \Delta'_{\mathbf{b}}, \varphi)$ of each buyer type $\mathbf{b}_i \in \Delta_{\mathbf{b}}$ as follows:

$$\phi(\mathbf{b}_i, \Delta'_b, \wp) = \min\{\phi'(\mathbf{b}_i, \Delta'_b, \wp), \text{RP}_{\mathcal{S}_{t+2}^1}\}$$

where $\phi'(\mathbf{b}_i, \Delta'_b, \wp)$ is the acceptance price of each buyer type $\mathbf{b}_i \in \Delta_b$ given the sequential equilibrium \wp :

$$\phi'(\mathbf{b}_i, \Delta'_b, \wp) = \begin{cases} \rho(\mathbf{b}_i, t, EBO(\mathbf{b}_i, \wp)) & \text{if } \mathbf{b}_i \in \Delta'_b \\ \min_{\mathbf{b}_j \in \Delta'_b} \rho(\mathbf{b}_i, t, EBO(\mathbf{b}_j, \wp)) & \text{otherwise} \end{cases}$$

Seller \mathcal{S}_{t+2}^1 's expected utility of making an offer x given the sequential equilibrium \wp is defined as

$$EU_s(x, t) = \sum_{\mathbf{b}_i \in \Delta_b} \omega_{\mathbf{b}_i}(\Delta_b) EU_s(x, t, \mathbf{b}_i)$$

where $EU_s(x, t, \mathbf{b}_i)$ is seller \mathcal{S}_{t+2}^1 's utility if the buyer is of type \mathbf{b}_i , which is defined as

$$\begin{cases} U_s(x, t+1) & \text{if } x \leq \phi(\mathbf{b}_i, \Delta'_b, \wp) \\ U_s(EBO(\mathbf{b}_i, \wp)) & \text{if } x > \phi(\mathbf{b}_i, \Delta'_b, \wp) \text{ and } \mathbf{b}_i \in \Delta'_b \\ U_s(\min_{\mathbf{b}_j \in \Delta'_b} \rho(\mathbf{b}_i, t, EBO(\mathbf{b}_j, \wp)), t+1) & \text{otherwise} \end{cases}$$

It is easy to see that the optimal offer the seller \mathcal{S}_{t+2}^1 should be either one buyer type's acceptance price or a price that will be rejected by all buyer types (i.e., ϖ). If the seller \mathcal{S}_{t+2}^1 's optimal offer x satisfies the first two equilibrium existence conditions, there is a sequential equilibrium in which the seller \mathcal{S}_{t+2}^1 offers price x and buyer types Δ'_b will reject the offer with the sequential equilibrium \wp . Otherwise there is no sequential equilibrium given this reject update rule and the continuation game equilibrium \wp .

In addition to the above reject update rules under which at least one buyer type will choose to reject the offer, the seller \mathcal{S}_{t+2}^1 can also make an offer such that it is all buyer types' equilibrium strategies to accept the offer. It is easy to see that the highest offer that will be accepted by all buyer types in equilibrium is $x = \min\{\text{RP}_{\mathcal{S}_{t+2}^2}, \min_{\mathbf{b}_i \in \Delta_{\mathbf{b}}}(x_{\mathbf{b}_i}^*(t+1)|\mathbf{b}_h(\Delta_{\mathbf{b}}))_{\leftarrow \mathbf{b}_i}\}$ since if a seller offers a price larger than x , at least one buyer type has an incentive to deviate from accepting the offer. If the buyer rejects x , the seller will update its belief to $\mathbf{b}_h(\Delta_{\mathbf{b}})$. The acceptance price of buyer type \mathbf{b}_i for this reject update rule is thus $(x_{\mathbf{b}_i}^*(t+1)|\mathbf{b}_h(\Delta_{\mathbf{b}}))_{\leftarrow \mathbf{b}_i}$. If the optimal offer of the seller in this case is not acceptable to all the buyer types (i.e., the optimal offer is not $\min_{\mathbf{b}_i \in \Delta_{\mathbf{b}}}(x_{\mathbf{b}_i}^*(t+1)|\mathbf{b}_h(\Delta_{\mathbf{b}}))_{\leftarrow \mathbf{b}_i}$), there is no sequential equilibrium for this null reject update rule. Otherwise, there is a sequential equilibrium in which the seller will make an offer which will be accepted by all buyer types.

4.4 Summary

This chapter analyzes agents' strategic behavior in concurrent one-to-many negotiation and many-to-many negotiation when agents follow the alternating-offers protocol. The analysis can provide insights and suggestions for designing negotiation agents in practical electronic marketplaces in which agents are involved in many-to-many negotiations. The contributions of this chapter can be summarized as follows:

- We extend the alternating-offers protocol to handle multiple trading opportunities and market competition. We provide an algorithm based on backward induction to compute the subgame perfect equilibrium of concurrent one-to-many negotiation and many-to-many negotiation. We observe that agents' bargaining power are affected by the proposing ordering and market competition.
- For the complete information setting, we show that the computational complexity when there are many buyers and many sellers in our protocol lineally

increases with the number of buyers and sellers. We find that for a large subset of the space of the parameters, agents' equilibrium strategies depend on the values of a narrow number of parameters. The computation of the equilibrium for realistic ranges of the parameters in one-to-many settings reduces to the computation of the equilibrium either in one-to-one settings with uncertainty or in one-to-many settings without uncertainty. We also compare the efficiency of the negotiation mechanism with that of some other mechanisms like VCG auction.

- We provide an algorithm to find a pure strategy sequential equilibrium in one-to-many negotiation where there is uncertainty regarding the reserve price of one agent. Our algorithm combines together game theoretic analysis with state space search techniques and it is sound and complete.

The assumptions made in this chapter are not more restrictive than related work in the literature. The assumption of the existence of deadline and reserve price in bargaining is widely used in the literature (e.g., [57, 61, 102, 112]). Computing agents' equilibrium strategies in incomplete information bargaining is extremely difficult and most related work only considers one type of uncertainty. For instance, Rubinstein [112] considered bilateral bargaining with uncertainty over two possible discount factors. Gatti *et al.* [61] analyzed bilateral bargaining with one-sided uncertain deadlines. In this chapter, we consider the uncertain information about the reserve price of an agent while assuming complete information about other negotiation parameters. As in most related work, we consider the negotiation over a single issue, price of a good. However, our analysis can be easily extended to the multi-attribute negotiations in which the attributes are negotiated simultaneously [61]. The next chapters of this thesis will consider more realistic negotiation problems and will present heuristic based negotiation strategies and also protocols.

CHAPTER 5

STRATEGIC AGENTS FOR MULTI-RESOURCE NEGOTIATION

This chapter investigates automated negotiation in resource allocation among resource providers (sellers) and consumers (buyers), where consumer agents may require multiple resources to successfully complete their tasks. Therefore, consumer agents may need to engage in multiple negotiations. If the multiple negotiations are not all successful, consumers gain nothing. This is a simple form of multi-linked negotiation where the resources are interrelated in the sense that, from the perspective of the overall negotiation, resources are dependent as an agent's utility from the overall negotiation depends on obtaining overall agreements on all the resources. This chapter presents the design and implementation of agents that concurrently negotiate with other entities for acquiring multiple resources.

5.1 Background

In electronic commerce markets where selfish agents behave individually, agents often have to acquire multiple resources in order to accomplish a high level task with each resource acquisition requiring negotiations with multiple resource providers. Such scenarios widely exist in practical applications. For example, a complex task may need several robots to work together and the absence of any robot results in the failure of the task. This is a simple form of multi-linked negotiation where the resources are independent but are interrelated. Resources are independent in the sense that there is no dependence between different resources, i.e., acquiring one resource

doesn't constrain how the other resources are acquired. However, from the perspective of the overall negotiation, resources are dependent as an agent's utility from the overall negotiation depends on obtaining overall agreements on all the resources. The negotiation problem in this chapter has the following three features:

1. When acquiring multiple resources, a consumer agent only knows the reserve price available for the entire set of resources, i.e., the highest price the agent can pay for all the resources, rather than the reserve price of each separate resource. In practice, given a plan and its resource requirements, an agent can easily decide the reserve price for all the resources in that plan based on the overall worth of the task. However, it is difficult (even impossible) for a resource consumer to understand how to set the reserve price for each separate resource. In fact, we show experimentally that it is undesirable to set a fixed reserve price for an individual resource prior to beginning negotiations.
2. Agents can decommit from tentative agreements at the cost of paying a penalty. Decommitment allows agents to profitably accommodate new tasks arriving or new negotiation events. If these events make some existing contracts less profitable or infeasible for an agent, that agent can decommit from those contracts [121].
3. Negotiation agents are assumed to have incomplete information about other agents, for example, a buyer agent knows the distribution of the reserve price of a seller agent and the number of trading competitors. However, an agent's negotiation status (the set of proposals it has received) and negotiation strategy are its private information. For strategic reasons, a negotiation agent won't disclose such information during negotiation. During negotiation, negotiation agents can quit negotiation at any time, even without notifying their trading

partners. When a buyer acquires multiple resources, it concurrently negotiates with sellers to reach agreements for all the resources.

Currently, there are limited techniques based on auctions or independent negotiations over single resources for performing the assembly of multiple resources required by a task. A centralized approach such as reverse combinatorial auctions [42, 99] requires a controlling agent (the auctioneer) for determining which agents receive which resources based on the bids submitted by individual agents. However, the auctioneer may face significant computational overload due to a large number of bids with complex structure. Assume that each buyer runs a reverse combinatorial auction, each seller may participate in multiple auctions as there are multiple buyers requiring its resource. It's difficult for each seller to derive its optimal bids for all the concurrent auctions. An alternative approach is that each buyer (seller) submits its resource requirement (supply) to a super agent and the super agent runs auctions for all the buyers (sellers). However, it may be difficult to find such an auctioneer agent that selfish agents can trust and can comply with the decisions made by the auctioneer. Moreover, in dynamic environments that resource supply and demand arrive randomly, it is very difficult for the auctioneer to decide optimally when to run auctions. In our distributed approach, allocations emerge as the result of a sequence of distributed negotiations and each selfish agent acts on behalf of itself. An agent can negotiate with other agents when needed. The distributed model is also more suitable for the situation when the needed resources are from multiple electronic marketplaces (i.e., no centralized auction is possible), and more natural in cases where resources belong to different selfish agents and finding optimal allocations may be (computationally) infeasible. We feel it is key that the acquisition of multiple resources necessary is seen as an integrated process in which the results/status of any one negotiation affects all other negotiations.

Because resource providers and consumers may have different goals, preferences, interests, and policies, the problem of negotiating an optimal allocation of resources within a group of agents has been found to be intractable both in terms of the amount of computation [45] and communication needed [47]. The multi-resource negotiation studied in this chapter is even more complex due to the possibility of agents' decommitting from previously made agreements. An agent's bargaining position in each round is determined by many factors such as market competition, negotiation deadlines, current agreement set, trading partners' proposals, and market dynamics. During each round of negotiation, an agent has to make decisions on how to proceed with each negotiation thread and there are many possible choices for each decision based on a variety of factors. Thus, it is difficult to construct an integrated framework in which all these factors are optimized concurrently. Rather than explicitly modeling these inter-dependent factors and then determining each agent's best decisions by an intractable combined optimization, this work tries to connect those inter-dependent factors indirectly and develops a set of heuristics to approximate agents' decision making during negotiation. The distinguishing feature of negotiation agents in this chapter is their flexibility; they can adjust 1) the number of tentative agreements for each resource and 2) the amount of concession by reacting to i) changing market conditions, and ii) the current negotiation status of all concurrently negotiating threads. In our approach, agents utilize a time-dependent negotiation strategy in which the reserve price of each resource is dynamically determined by 1) the likelihood that negotiation will not be successful (*conflict probability*), 2) the *expected agreement price* of the resource, and 3) the *expected number of final agreements* given the set of tentative agreements made so far. The negotiation deadline of each resource is determined by both its scarcity and the overall deadline for the entire negotiation. A buyer agent can make more than one tentative agreement for each resource and the maximum

number of tentative agreements is constrained by the market situation in order to avoid the agent's making more agreements than necessary.

Our work here is connected to several lines of research in agent-mediated negotiation including multi-issue negotiation (e.g., [49, 50, 51, 80, 81, 82, 130, 133]), one-to-many negotiation [13, 14, 28, 96, 97, 106], negotiation strategies (e.g., [48, 78, 124, 125, 127]), and decommitment (e.g., [5, 98, 121]). This chapter presents the first design of negotiation agents in dynamic and uncertain environments in which 1) a consumer negotiates for multiple resources and its negotiation fails if it fails to get some resources, and 2) agents can choose to decommit from existing agreements within a fixed period. This research is intellectually challenging because of both the complex interactions among concurrent negotiations for multiple resources and the uncertainty associated with the outcome of these negotiations. This research provides a deep understanding of the influence of sophisticated negotiation mechanisms on individual agents' performance in dynamic environments, and hence contributes to the construction of effective problem-solving approaches in open environments. The proposed approach can be used for designing negotiation agents in many practical applications like service composition [105], grid resource management [126], and supply chain [138].

5.2 Negotiation mechanism

5.2.1 Assumptions

We make the following assumptions about agents' knowledge and strategies:

1) Agents have incomplete information about each other. The assumption of incomplete information is intuitive because in practice, agents have private information, and for strategic reasons, they do not reveal their strategies, constraints, or preferences. In [109, p.54], it was noted that the strategy of a trading agent corresponds to its internal program, and extracting the true internal decision process would be

difficult. Moreover, when selfish agents have competing interests, they may have incentive to deviate from protocols or to lie to other agents about their preferences. This chapter assumes that 1) agents know the number of trading partners and competitors and 2) the distributions of trading partners' reserve price. The assumption that the number of trading partners is known is less restrictive or similar to the assumptions in most related work (e.g., [50, 87, 96, 97, 98]). We consider both assumptions are realistic in practice. For example, consider the streaming processing system CLASP [27], in which each resource provider (consumer) always posts its resource supply (requirement). Further, the distribution of trading partners' reserve prices can be learned as a result of repeated interaction with agents in the marketplace. We explored the sensitivity of these assumptions in the experiment section.

2) A consumer agent negotiates over multiple resources in parallel and, for each resource, the agent concurrently negotiates with its trading partners. Given that the buyer doesn't know how to appropriately set the reserve price of each of its resources, one approach that requires no prior knowledge of the marketplace about current resource scarcity and expected competition of a specific resource is for a consumer to negotiate over all the resources in parallel. For each resource, there are multiple trading partners and the agent concurrently negotiates with all the trading partners. Therefore, each negotiation thread of one resource has multiple concurrently existing outside options. Generally, a buyer obtains more desirable negotiation outcomes when it negotiates concurrently with all the sellers in competitive situations in which there is information uncertainty and there is a deadline for the negotiation to complete [96, 97]. Additionally, inefficiency may arise in sequential negotiation when considering the overall time cost to complete all the necessary negotiations [51].




Figure 5.1. Buyer a 's multi-resource negotiation problem

5.2.2 The Negotiation Problem

All the analysis in this chapter is from the perspective of a randomly selected buyer a (see Figure 5.1). Let $\mathcal{I} = \{I_1, I_2, \dots, I_l\}$ be the set of resources needed by a and τ be a 's negotiation deadline. Let a negotiation period of a be denoted by t , $t \in \{0, 1, \dots, \tau - 1\}$. For resource I_j , a has a set \mathcal{TP}_j^t of trading partners (sellers) at round t . Also, a has a set \mathcal{CP}_j^t of trading competitors (buyers) for resource I_j at round t . $\phi_{a \rightarrow s}^t$ is the proposal of a to its trading partner $s \in \mathcal{TP}_j^t$ at round t . $\phi_{s \rightarrow a}^t$ is the proposal of seller agent s to a at round t . RP and IP are respectively, the reserve price (maximum amount of money a can spend) and the desirable price of a before negotiation begins, respectively. IP_j is a 's initial proposal price for resource I_j , i.e., $\phi_{a \rightarrow s}^0$, and it follows that $\sum_j IP_j = IP$. RP^t is a 's reserve price for all negotiating resources \mathcal{I}^t at round t . Once a tentative agreement (defined below) for I_j becomes a final agreement, a doesn't need further negotiation about I_j . Therefore, $\mathcal{I}^t \subseteq \mathcal{I}^{t-1} \subseteq \mathcal{I}$.

An agent can decommit from an agreement within λ rounds after the agreement has been made. Assume a makes an agreement Ag about resource I_j with agent s at round $\text{Tm}(Ag) = t$ and the agreement price is $\text{Prc}(Ag)$. Assume a decommits from the agreement Ag at round t' where $t' - \text{Tm}(Ag) \leq \lambda$. The penalty of the decommitment is defined by $\rho(\text{Prc}(Ag), t, t', \lambda)$. This chapter assumes that 1) penalty functions are nonnegative, continuous, and nondecreasing with time and agreement price, and 2) the maximum penalty is less than the agreement price. Therefore, if an agent makes unnecessary agreements for a resource, it will decommit from these unnecessary agreements. An example of such a penalty function is $0.1 \times \text{Prc}(Ag) \times ((t' - t)/\lambda)^\varsigma$ where $\varsigma > 0$.

Penalties could be different from one resource to another resource. If the two parties decommit at the same time, they don't need to pay a penalty to each other. An agreement made in the bargaining process is called a *tentative* agreement and it becomes a *final* agreement if neither party decommits from the agreement in the λ rounds after the agreement was made. Agent a needs to fulfill all its final agreements, i.e., a needs to pay for all final agreements, even through it needs only one final agreement for each resource. a tries to make agreements for all its resources and a gains nothing if it fails to make an agreement for any resource in I , no matter how many and how good the agreements for other resources are. In other words, a requires a set of resources and only receives a positive utility if it acquires all of them, and zero otherwise. This assumption makes sense in some practical domains like some supply chain or Grid applications where the failure of one step (or one sub-task) will result in the failure of the whole task. The utility function of a when a makes at least one final agreement for each resource is defined as:

$$u_a = RP - \sum_{I_j \in \mathcal{I}} \sum_{Ag \in \mathcal{FAG}_j^{\tau+\lambda}} \text{Prc}(Ag) + \sum_{t=0}^{\tau+\lambda} (\rho_{in}^t - \rho_{out}^t)$$

where $\tau + \lambda$ is the maximum period that a was involved in negotiation and decommitment, $\mathcal{FAG}_j^{\tau+\lambda}$ is the set of final agreements for resource I_j at $\tau + \lambda$, ρ_{out}^t is the penalty a pays to other agents at t when it decommits, and ρ_{in}^t is the payment of penalty a receives from other agents at t if they decommit.

If a fails to make a final agreement for at least one resource, a gains nothing and its utility is defined as:

$$u_a = - \sum_{I_j \in \mathcal{I}} \sum_{Ag \in \mathcal{FAG}_j^{\tau+\lambda}} \text{Prc}(Ag) + \sum_{t=0}^{\tau+\lambda} (\rho_{in}^t - \rho_{out}^t)$$

In this case, a does not get the value RP since its task cannot be completed and thus its utility may be negative. Its only “income” in this case is the penalty received from its trading partners.

5.2.3 The Negotiation Protocol

As agents can choose to decommit from agreements, negotiation consists of a *bargaining stage* and a *decommitment stage* for each negotiation thread. This work adopts the well known alternating offers protocol (see [111, p.100]) so that a pair of buyer and seller agents in a negotiation thread bargain by making proposals to each other. At each round, one agent makes a proposal first, then the other agent has three choices in the bargaining stage: 1) accept the proposal, 2) reject the proposal, or 3) make a counter proposal. For ease of analysis, this work assumes that buyers always propose first to sellers during negotiation. Many buyer-seller pairs can bargain simultaneously since each pair is in a negotiation thread. If the seller accepts the proposal of the buyer, negotiation terminates with a tentative agreement. If the seller rejects the proposal of the buyer, negotiation terminates with no agreement. If the seller makes a counter proposal, bargaining proceeds to another round and the buyer can accept the proposal, reject the proposal, or make a counter proposal. Bargaining between two agents terminates 1) when an agreement is reached or 2)

with a conflict (i.e., no agreement is made) when one of the two agents' deadline is reached or one agent quits the negotiation. After a tentative agreement is made, an agent has the opportunity to decommit from the agreement and the decommitting agent pays the penalty to the other party involved in the decommitted agreement.

5.2.4 The Negotiation Strategy

An agent's negotiation strategy is a function from the negotiation history to its actions at each negotiation round [109]. An agent a 's negotiation strategy can be represented as a sequence of functions $f_a = \{f_a^t\}_{t=0}^{\infty}$, where f_a^t is a 's strategy at round t . As the agent is negotiating for multiple resources and there are multiple negotiation threads for each resource, the agent's negotiation strategy f_a^t specifies for the agent what to do at round t for each of the active negotiation thread. For each trading partner s , the agent a has four choices: 1) accept the proposal by s , 2) reject the proposal by s , 3) make a counter proposal to s in the bargaining stage, or 4) decommit from the agreement between a and s in the decommitment stage.

A strategy profile $\mathcal{F} = (f_a, f_{\mathcal{TP}}, f_{\mathcal{CP}})$ is a collection of strategies, one for each agent, where $f_{\mathcal{TP}}$ and $f_{\mathcal{CP}}$ are the strategies for a 's trading partners and trading competitors, respectively. Let $\mathfrak{S} : \mathcal{F} \rightarrow \mathcal{O}$ be a social choice function which determines the negotiation result given the negotiation strategies \mathcal{F} of all the agents. Given the strategy profile of all the agents, game theory has been widely applied in analyzing the equilibria of bargaining models (e.g., Nash equilibria, Sub-game perfect equilibria, Sequential equilibria) [100]. The analytic complexity of equilibrium analysis increases rapidly when more elements (e.g., deadline, outside options, bargaining costs, market competition) and more agents are included in the model. As a result, in most models, only one or two elements are considered. For example, Rubinstein [111] studies a two-player sequential bargaining game in which bargaining cost is considered. The latest advance in computing sequential equilibrium strategies only considers a bilateral bar-

gaining model in which one agent has incomplete information about the deadline of the other agent [61]. We take a set of elements into account, for example, deadline, outside option, market competition, multiple resources, and decommitment. In addition, we are not assuming that agents have complete information about the factors considered in our framework, which makes agents' reasoning even more difficult. Therefore, we feel that it is impractical to formally model the complex interaction that occurs between the bargaining and decommitment nor the interaction among multiple resources in the framework.

If we assume that each agent has information, which could be a probabilistic distribution, about other agents' strategies (i.e., $f_{\mathcal{TP}}$ and $f_{\mathcal{CP}}$), the optimization problem of agent a is to find the optimal negotiation strategy f_a^* from the set \mathcal{F}_a of possible negotiation strategies:

$$f_a^* = \operatorname{argmax}_{f_a \in \mathcal{F}_a} u_a(\mathfrak{S}(f_a, f_{\mathcal{TP}}, f_{\mathcal{CP}}))$$

where $u_a(\mathfrak{S}(f_a, f_{\mathcal{TP}}, f_{\mathcal{CP}}))$ is a 's utility of the negotiation result $\mathfrak{S}(f_a, f_{\mathcal{TP}}, f_{\mathcal{CP}})$. Agent a 's optimization problem at each negotiation round t can be formulated as a Markov Decision Process (MDP) $\langle S, A, P, R \rangle$ where the state set S can be characterized by the market situation (e.g., the number of buyers or sellers, the agreement set of each buyer or seller), action set A consists of all the actions each agent can choose (e.g., a counter-proposal including the price, or decommitment decision), transition function P is determined by agents' negotiation strategies and the change of market with time, reward function R is based on the utility each agent can gain from a specific state. As the action space A is infinite, solving the MDP problem could be computationally intractable [22]. Moreover, as stated before, it's impractical to assume that agents have information about other agents' negotiation strategies. For strategy or privacy reasons, an agent is unwilling to broadcast its decisions.

Given that 1) it's hard (even impossible) to compute agents' equilibrium strategies, and 2) it's not appropriate to assume that a knows other agents' negotiation strategies, this chapter presents a set of heuristics for agents to make negotiation decisions at each negotiation round. The set of heuristics consider many relevant issues such as the risk that their negotiation partners may decommit (and therefore the fact that ideally a buyer needs to secure more than one agreement for any given resource), the competition that buyers face from other buyers, uncertainty about the reserve prices of their trading partners, multiple opportunities of reaching an agreement, the set of available tentative agreements, deadline, and negotiation history.

5.3 Heuristics based Strategies

Agent a has l resources to acquire, and for each resource, a conducts multi-threaded negotiation with a set of trading partners. For each negotiation thread associated with the acquisition of a resource, a needs to decide 1) what is its proposal during the bargaining stage and 2) when and whether to decommit from an agreement in the decommitment stage.

5.3.1 An overview of negotiation strategies

Algorithm 3 gives an overview of a 's strategy during the bargaining stage and the decommitment stage.

At round $t = 0$, a needs to make an initial proposal IP_j to each trading partner s . During each later round ($t > 0$), a will always first update its information structures (see Algorithm 4). First, if another agent decommits from an agreement, then remove the agreement from the tentative agreement set. Second, if another agent sends a message indicating rejection of the current proposal, the corresponding negotiation thread terminates. If another agent accepts a proposal, then add the agreement to the tentative agreement set. If one tentative agreement becomes a final agreement (no

Table 5.1. Symbols used in this chapter

τ	deadline of agent a
I_j	resource j
RP	reserve price for all resources
IP	desirable price for all resources
IP_j	initial proposal for resource I_j
τ_j^t	deadline of agent a for resource I_j at round t
\mathcal{I}^t	the set of resources at round t
\mathcal{TP}_j^t	the set of partners (sellers) about I_j at round t
\mathcal{CP}_j^t	the set of competitors (buyers) about I_j at round t
$\phi_{a \rightarrow s}^t$	a 's proposal to s at round t
\mathcal{P}_j^t	a 's trading partners' proposals about I_j at round t
RP^t	a 's reserve price for all negotiating resources at round t
RP_j^t	a 's reserve price for resource I_j at round t
\mathcal{TAG}_j^t	a 's set of tentative agreements for resource I_j at round t
\mathcal{FAG}_j^t	a 's set of final agreements for resource I_j at round t
$\text{Prc}(Ag)$	price of the agreement Ag
$\text{Tm}(Ag)$	time when the agreement Ag was made
ρ_{out}^t	the penalty a pays to other agents at round t
ρ_{in}^t	the payment of penalty a receives at round t
C_j^t	the scarcity of resource I_j at t
RC_j^t	the relative scarcity of resource I_j at t
δ_j^t	the concession rate with respect to resource I_j at round t
χ_j^t	the conflict probability of the negotiation for I_j at t
ϖ_j^t	the expected agreement price of resource I_j at t
$\omega_s^t(Ag)$	the probability of s 's decommitting from Ag at t
$\varphi(\mathcal{TAG}_j^t)$	the expected number of final agreements given \mathcal{TAG}_j^t
$\gamma(\mathcal{TAG}_k^t)$	model how $\varphi(\mathcal{TAG}_j^t)$ affects the offering price

decommitment allowed) for the resource I_j as the negotiation moves to a new round, then a will decommit from all tentative agreements about I_j , stop all negotiation threads for I_j , and remove I_j from \mathcal{I}^t .¹

Next a computes the negotiation deadline τ_j^t for each resource $I_j \in \mathcal{I}^t$ (Section 5.3.2) and generates a proposal $\phi_{a \rightarrow s}^t$ to each trading partner $s \in \mathcal{TP}_j^t$ (Section 5.3.3). If $\phi_{a \rightarrow s}^t < \phi_{s \rightarrow a}^{t-1}$ (i.e., s 's last proposal is not acceptable), then a sends the proposal $\phi_{a \rightarrow s}^t$ to s directly. Otherwise, it adds $\langle \phi_{s \rightarrow a}^{t-1}, t \rangle$ into tentative agreement set \mathcal{TAG}_j^t .

For resource I_j , a checks whether the current set of agreements are sufficient. If the current set of agreements is more than needed, a recursively removes agreements from the tentative agreement set (Section 5.3.4). Assume that Ag needs to be removed and the trading partner in the agreement Ag is seller s . If $Ag \in \mathcal{TAG}_j^{t-1}$, then a decommits from the agreement. If Ag is not in \mathcal{TAG}_j^{t-1} , the agreement Ag has been just added to \mathcal{TAG}_j^t by a at time t but the seller involved in the agreement hasn't received the "accept" message from a . Although a doesn't intend to make the agreement Ag and a can quit the negotiation with s , it's better for a to continue the negotiation with s and try to get better agreements than an agreement in the current tentative agreement set \mathcal{TAG}_j^t . Therefore, a removes Ag from \mathcal{TAG}_j^t and sends s a proposal with lower price than the price in the agreement Ag . Finally, if an agreement Ag is contained in \mathcal{TAG}_j^t but is not in \mathcal{TAG}_j^{t-1} , then a sends an *accept* proposal to the corresponding seller involved in the agreement Ag .

¹The only additional value that can be achieved by keeping alive any future negotiation is the possibility that a trading partner is likely to decommit. In this case, it would be profitable to delay decommitment and thus the agent does not need to pay the decommitment penalty but receives the penalty from its trading partner. However, since each buyer does not know whether a seller will decommit from an agreement and the penalty increases with time, the buyer may have to pay a higher penalty if it has to decommit before the unnecessary tentative agreement becomes a final agreement. We evaluated the benefit of delaying decommitment through experimentation and found that delaying decommitment did not increase the buyer's average utility. In the current implementation, we do not take this into account.

Algorithm 3: Negotiation Strategy of Agent a

Data Structure: Tentative agreement set \mathcal{TAG}_j^t , final agreement set \mathcal{FAG}_j^t , sellers' proposal set for each resource I_j at round t .

Output: Final agreement set \mathcal{FAG}_j^t for each I_j

- 1: Initial proposing: Let $t = 0$ and propose IP_j to every trading partner s about I_j .
 - 2: **repeat**
 - 3: $t++$;
 - 4: $\mathcal{I}^t = \mathcal{I}^{t-1}$;
 - 5: $\mathcal{TAG}_j^t = \mathcal{TAG}_j^{t-1}$, $\mathcal{FAG}_j^t = \mathcal{FAG}_j^{t-1}$ for $I_j \in \mathcal{I}^t$;
 - 6: **Step 1: initialization (Algorithm 2)**
 - 7: **Step 2: deadline calculation (Section 5.3.2)**
 - 8: **Step 3: proposal generation (Section 5.3.3)**
 - 9: **Step 4: meet the agreement number constraint (Section 5.3.4)**
 - 10: **Step 5: send left proposals**
 - 11: **until** 1) $t \geq \tau + \lambda$, or 2) $|\mathcal{FAG}_j^t| > 0$ for each I_j , or 3) $|\mathcal{TAG}_j^t| = 0$ for some I_j at $t \geq \tau_j^t$
-

The overall negotiation process will terminate if 1) the deadline is reached, or 2) a makes a final agreement for each resource I_j , or 3) $|\mathcal{TAG}_j^t| = 0$ for some I_j at $t \geq \tau_j^t$, which means it no longer makes any sense for a to make any other agreements.

This work assumes that a buyer agent always offers the same price to all trading partners of one resource. Formal analysis of concurrent negotiation [7] suggests that it is an agent's dominant strategy to make the same offer to all trading partners. While this chapter considers more complex negotiation, it is still intuitive to not make price discrimination proposals for the same resource. While making an offer, a buyer hopes that the offer would be accepted. If there are two offers which have the same probability of being accepted, the buyer will choose the offer with the lower price.

5.3.2 Different deadlines for different resources

The number of buyers and sellers for different resources varies. A resource is easy to obtain if the number of sellers is much larger than the number of buyers. In contrast, if there are more buyers and less sellers, the resource is relatively difficult to

Algorithm 4: Initialization

```
1: for each  $I_j \in \mathcal{I}^t$  do
2:   for each  $s \in \mathcal{TP}_j^{t-1}$  do
3:     if  $\phi_{s \rightarrow a}^t = \text{“decommit from } Ag\text{”}$  then
4:       remove  $Ag$  from  $\mathcal{TAG}_j^t$ 
5:     else
6:       if  $\phi_{s \rightarrow a}^t = \text{“reject”}$  then
7:         remove  $s$  from  $\mathcal{TP}_j^t$ 
8:       end if
9:     else
10:      if  $\phi_{s \rightarrow a}^t = \text{“accept”}$  then
11:        add  $\langle \phi_{s \rightarrow a}^{t-1}, t \rangle$  into  $\mathcal{TAG}_j^t$ 
12:      end if
13:    end if
14:  end for
15:  for each  $Ag \in \mathcal{TAG}_j^t$  do
16:    if  $t - \text{Tm}(Ag) > \lambda$  then
17:      remove  $Ag$  from  $\mathcal{TAG}_j^t$  and add it to  $\mathcal{FAG}_j^t$ 
18:    end if
19:  end for
20:  if  $|\mathcal{FAG}_j^t| > 0$  then
21:    decommit from all agreements in  $\mathcal{TAG}_j^t$ , stop all negotiation threads for  $I_j$ ,
    and remove  $I_j$  from  $\mathcal{I}^t$ .
22:  end if
23: end for
```

obtain since the resource seems “scarce” in terms of the ratio of supply to demand. The intuition behind using different negotiation deadlines for different resources is based on the following scenario: a makes an agreement about a scarce resource I_j before the deadline approaches. However, the other party involved in the agreement later decommits from the agreement. Then, the overall negotiation fails as it’s difficult for agent a to get another agreement for the scarce resource I_j in the remaining time and thus a needs to pay the penalty for its other agreements. To decrease the possibility of this situation happening, we can reduce the deadlines of scarce resources to increase the likelihood that we have a final agreement for those resources in place before the overall negotiation deadline. In other words, we would like to quickly secure one final agreement for a scarce resource. On one hand, by decreasing one resource’s

artificial deadline, a is inclined to make larger concessions to its trading partners and thus its probability of making a final agreement for the resource increases. On the other hand, if it's difficult for a to make a final agreement for one resource, a can know this earlier. Thus a can pay less decommitment penalties by decommitting from agreements earlier as penalties increase with time. However, the determination of this virtual deadline for scarce resources is a dynamic process which can either decrease or increase the deadline as conditions change in the future.

The scarcity of a resource I_j is evaluated based on the competition situation of the negotiation over resource I_j . A negotiator's bargaining "power" is affected by the number of competitors and trading alternatives. Multiple options give a negotiator more "power" since the negotiating party needs not pursue the negotiation with any sense of urgency. The competition situation of an agent is determined by the probability that it is considered as the most preferred trading partner [127]. An agent's preferred trading partner refers to the one who makes the best proposal to the agent. a has \mathcal{CP}_j^t competitors and \mathcal{TP}_j^t partners. While it's impossible for a to compute exactly the probability that it is considered as the most preferred trading partner since a doesn't know other agents' negotiation strategies, the probability can be approximated in the following way. The probability that a is not the most preferred trading partner of any trading partner is $\mathcal{CP}_j^t/(\mathcal{CP}_j^t + 1)$. The probability of the agent a not being the most preferred trading partner of all the trading partners is approximated by

$$C_j^t = \left(\frac{\mathcal{CP}_j^t}{\mathcal{CP}_j^t + 1} \right)^{\mathcal{TP}_j^t}$$

C_j^t measures the scarcity of resource I_j at t . With more trading partners, it is relatively less difficult to acquire the resource and C_j^t will decrease. With more trading competitors, it is relatively more difficult to acquire the resource and C_j^t will increase.

If resource I_j is scarce and the other resources are not scarce, it's reasonable to decrease I_j 's deadline in order to decrease the probability that the overall negotiation

fails due to the failure of the negotiation about resource I_j . However, if all the desired resources are scarce, it may not be necessary to decrease the deadline of all the resources. In other words, whether to decrease the deadline of the resource I_j may not depend on the absolute scarcity of the resource, but rather its “relative scarcity”. The relative scarcity of the resource I_j is defined as the ratio of the I_j ’s scarcity measure to the harmonic mean of the scarcity measure of all the resources:

$$RC_j^t = \frac{C_j^t}{\frac{|\mathcal{I}^t|}{\sum_{I_k \in \mathcal{I}^t} \frac{1}{C_k^t}}} = \frac{C_j^t \sum_{I_k \in \mathcal{I}^t} \frac{1}{C_k^t}}{|\mathcal{I}^t|}$$

Using harmonic mean, the scarcer resource dominates the deadline calculation, which is close to the practice. Given the relative scarcity of each resource $I_j \in \mathcal{I}^t$, the deadline of resource I_j at time t is given as follows

$$\tau_j^t = \begin{cases} \tau & \text{if } RC_j^t < 1 \\ (RC_j^t)^\rho \tau & \text{if } RC_j^t \geq 1 \end{cases}$$

where $\rho < 0$. If the resource I_j is not scarce as compared with most resources, the deadline for resource I_j will be the deadline of the overall negotiation. Otherwise, i.e., $RC_j^t \geq 1$, its deadline τ_j^t is smaller than τ as $(RC_j^t)^\rho < 1$, and it can be found that τ_j^t will decrease with the increase of RC_j^t . That is, a relatively scarcer resource will have a shorter deadline.

5.3.3 Generating proposals

Since bargaining is fundamentally time-dependent [78, 48], agents utilize a time-dependent strategy when making concessions. Assume that a is negotiating with s about resource I_j . Then, a ’s proposal to s at round t is given by:

$$\phi_{a \rightarrow s}^t = IP_j + (RP_j^t - IP_j)\delta_j^t$$

where RP_j^t is agent a 's current reserve price of resource I_j at round t and δ_j^t is agent a 's concession rate with respect to resource I_j at round t , which is given by

$$\delta_j^t = T(t, \tau_j^t, \varepsilon) = (t/\tau_j^t)^\varepsilon$$

With infinitely many values of ε , there are infinitely many possible strategies in making concessions with respect to the remaining time. However, they can be classified into: 1) *Linear*: $\varepsilon = 1$, 2) *Conciliatory*: $0 < \varepsilon < 1$, and 3) *Conservative*: $\varepsilon > 1$ [127]. ε reflects an agent's mental state about its eagerness for finishing the negotiation earlier [78, 48]. Before making proposals, a needs to decide its reserve price RP_j^t . To calculate RP_j^t , we consider three factors: 1) the conflict probability χ_j^t which measures the aspiration level of the current negotiation for resource I_j , 2) the expected agreement price ϖ_j^t of resource I_j , and 3) the expected number $\varphi(\mathcal{TAG}_j^t)$ of final agreements based on the estimation of the decommitment probabilities of the current tentative agreement set. Function $\gamma(\mathcal{TAG}_k^t)$ is used to model the effect of the expected number $\varphi(\mathcal{TAG}_j^t)$ of final agreements.

RP_j^t is defined as:

$$RP_j^t = RP^t \frac{\chi_j^t \varpi_j^t \gamma(\mathcal{TAG}_j^t)}{\sum_{I_k \in \mathcal{I}^t} \chi_k^t \varpi_k^t \gamma(\mathcal{TAG}_k^t)}$$

where $RP^t = RP - \sum_{I_j \in \mathcal{I}} \sum_{Ag \in \mathcal{FAG}_j^t} \text{Prc}(Ag) + \sum_{t=0}^{t-1} (\rho_{in}^t - \rho_{out}^t)$ is agent a 's reserve price for all resources at round t , i.e., the maximum amount of money that it can spend to acquire all the remaining resources. We can see that the reserve price RP_j^t increases with the increase of the conflict probability χ_j^t and expected agreement price ϖ_j^t . If the current negotiation for resource I_j seems difficult, a needs to set a higher reserve price for resource I_j . Similarly, a needs to set a higher reserve price for resource I_j if the expected agreement price for resource I_j is high. Later we will show that $\gamma(\mathcal{TAG}_j^t)$ decreases with the increase of $\varphi(\mathcal{TAG}_j^t)$. Thus, the reserve price RP_j^t

decreases with the increase of the expected number $\varphi(\mathcal{TAG}_j^t)$ of final agreements, which is intuitive as buyers don't need to set a higher reserve price for a resource I_j when a has already made enough tentative agreements for I_j .

Conflict probability χ_j^t : Suppose that at round t , a 's last proposal $\phi_{a \rightarrow s}^{t-1}$ generates a utility of v_a for itself and v_s for s , and its trading partner s 's proposal $\phi_{s \rightarrow a}^{t-1}$ generates a utility of w_s for itself and w_a for a . Since a and s are utility maximizing agents, $v_a > w_a$ and $v_s < w_s$. If a accepts s 's last proposal, then it will obtain w_a with certainty. If a insists on its last proposal and 1) s accepts it, a obtains v_a and 2) s does not accept it, a may be subjected to a conflict utility c_a . c_a is the worst possible utility for a (i.e., a 's utility in the absence of an agreement with s). If s does not accept a 's last proposal, a may ultimately have to settle with lower utilities (the lowest possible being the conflict utility), if there are changes in the market situation in subsequent cycles. For instance, a may face more competitions in the next or subsequent cycles and may have to ultimately accept a utility that is lower than w_a (even c_a). If the subjective probability of obtaining c_a is p_c (conflict probability) and the probability that a achieving v_a is $1 - p_c$, and if a insists on holding its last proposal, a will obtain a utility of $(1 - p_c)v_a + p_c c_a$. Hence, a will find that it is advantageous to insist on its last proposal only if

$$(1 - p_c)v_a + p_c c_a \geq w_a$$

i.e., $p_c \leq (v_a - w_a)/(v_a - c_a)$ [124, 125, 127]. The maximum value of $p_c = (v_a - w_a)/(v_a - c_a)$ is the highest probability of a conflict that a may encounter in which $v_a = RP_j^t - \phi_{a \rightarrow s}^{t-1}$ and $w_a = RP_j^t - \phi_{s \rightarrow a}^{t-1}$. p_c is a ratio of two utility differences. While $v_a - w_a$ measures the cost of accepting the trading agent's last proposal, $v_a - c_a$ measures the cost of provoking a conflict. $v_a - c_a$ represents the range of possible values of utilities between the best case utility and the worst case (conflict) utility.

If there is no tentative agreement for resource I_j , i.e., $|\mathcal{TAG}_j^t| = 0$, the worst case utility c_a is 0. If $|\mathcal{TAG}_j^t| > 0$, a can use one of its tentative agreements as the finally agreement and c_a is defined as

$$\max_{Ag \in \mathcal{TAG}_j^t} (RP_j^t - \text{Prc}(Ag) - \text{Pnt}(\mathcal{TAG}_j^t - Ag, t, \lambda))$$

where $\text{Pnt}(\mathcal{TAG}, t, \lambda)$ is an estimation of the penalty a needs to pay while decommitting from the set of agreements \mathcal{TAG} . $\text{Pnt}(\mathcal{TAG}, t, \lambda)$ is defined as

$$\sum_{Ag \in \mathcal{TAG}} \frac{\sum_{t'=t}^{\text{Tm}(Ag)+\lambda} \rho(\text{Prc}(Ag), \text{Tm}(Ag), t', \lambda)}{\text{Tm}(Ag) + \lambda - t + 1}$$

in which any agreement $Ag \in \mathcal{TAG}$ can be decommitted at any time before the decommitment stage expires.

Aggregated Probability of Conflict: Let p_c^i be the conflict probability of a with any of its trading partner s and w_a^i be a 's utility if it accepts s 's proposal, then the aggregated conflict probability of a with all of its trading partners about I_j is given as follows [124, 125, 127]:

$$\chi_j^t = \prod_{i=1}^{|\mathcal{TP}_j^t|} p_c^i = \prod_{i=1}^{|\mathcal{TP}_j^t|} \frac{v_a - w_a^i}{v_a - c_a} = \frac{\prod_{i=1}^{|\mathcal{TP}_j^t|} (v_a - w_a^i)}{(v_a - c_a)^{|\mathcal{TP}_j^t|}}$$

Expected agreement price ϖ_j^t : Different resources have different ranges of agreement prices. For example, you may need to spend \$20,000 for a car but only need \$500 for a bike. Therefore, it's necessary to consider a resource's expected agreement price ϖ_j^t while determining the reserve price of the resource. ϖ_j^t is computed based on agent a 's estimation of the reservation price of a trading partner. The estimation is characterized by a probability distribution $F_s(\cdot)$, where $F_s(y)$ denotes the probability

that the reservation price of a trading partner s is no greater than y . $F_s(y)$ is identical and independent across all sellers.² This probability distribution is the prior belief of the buyer. For simplicity, let $F_j(y) = F_s(y)$ denote the probability that the reservation price of any trading partner $s \in \mathcal{TP}_j^t$ is no greater than y . The probability density function of $F_j(y)$ is denoted by $f_j(y)$. The desirable price IP_j for resource I_j is simply computed by considering sellers' reserve price for resource I_j : $IP_j = \int_{-\infty}^{\infty} f_j(y)ydy$.

Let $F_j^k(y)$ be the probability distributions of the k^{th} highest maximum reserve price. The probability density function of $F_j^k(y)$ is denoted by $f_j^k(y)$. $F_j^1(y)$ is equal to the product of the probabilities that the maximum reserve price is less than or equal to y in each thread. $F_j^2(y)$ is equal to $F_j^1(y)$ plus the probability that the highest maximum reserve price is greater than y , and the second highest maximum reserve price is less than or equal to y . These probabilities can be calculated by the following formulas:

$$\begin{aligned}
F_j^1(y) &= (F_j(y))^{|T\mathcal{P}_j^t|} \\
F_j^2(y) &= F_j^1(y) + C_{|T\mathcal{P}_j^t|}^1 (1 - F_j(y))^{2-1} (F_j(y))^{|T\mathcal{P}_j^t|-1} \\
F_j^k(y) &= F_j^{k-1}(y) + C_{|T\mathcal{P}_j^t|}^{k-1} (1 - F_j(y))^{k-1} (F_j(y))^{|T\mathcal{P}_j^t|-k+1}
\end{aligned}$$

The corresponding probability density functions are:

²Our model can also be extended to allow $F_s(y)$ to be different for different trading partners.

$$\begin{aligned}
f_j^1(y) &= |\mathcal{TP}_j^t| (F_j(y))^{|\mathcal{TP}_j^t|-1} \\
f_j^2(y) &= f_j^1(y) - C_{|\mathcal{TP}_j^t|}^1 f_j(y) (F_j(y))^{|\mathcal{TP}_j^t|-1} + C_{|\mathcal{TP}_j^t|}^1 \\
&\quad (|\mathcal{TP}_j^t| - 1) f_j(y) (1 - F_j(y))^{2-1} (F_j(y))^{|\mathcal{TP}_j^t|-2} \\
f_j^k(y) &= f_j^{k-1}(y) - C_{|\mathcal{TP}_j^t|}^{k-1} (k-1) f_j(y) \\
&\quad (1 - F_j(y))^{k-2} (F_j(y))^{|\mathcal{TP}_j^t|-k+1} \\
&\quad + C_{|\mathcal{TP}_j^t|}^{k-1} f_j(y) (1 - F_j(y))^{k-1} (F_j(y))^{|\mathcal{TP}_j^t|-k+1}
\end{aligned}$$

We provide a heuristic approach to estimate the expected agreement price for resource I_j . When the number of trading partners is less than the number of trading competitors, the agreement price follows the highest maximum reserve price distribution. Otherwise, the agreement price follows a lower reserve price distribution. This is also the case with less trading competitors. The intuition behind the heuristic is as follows. Consider the single-shot negotiation between buyers and sellers in which buyers make offers first and then sellers decide whether to accept or not. If there is no competitors, the equilibrium offer of the buyer a is sellers' lowest reserve price. If there is one competitor, the equilibrium offer of the buyer a is the second lowest reserve price. In the same way, if there are $|\mathcal{CP}_j^t|$ competitors, the equilibrium offer is the $(|\mathcal{CP}_j^t| + 1)^{th}$ lowest reserve price, i.e., $(|\mathcal{TP}_j^t| - |\mathcal{CP}_j^t|)^{th}$ highest reserve price. Since in our model buyers don't know sellers' exact reserve prices, distributions are used instead. Formally, ϖ_j^t is given as follows:

$$\varpi_j^t = \begin{cases} \int_{-\infty}^{\infty} f_j^{|\mathcal{TP}_j^t|-|\mathcal{CP}_j^t|}(y) y dy & \text{if } |\mathcal{TP}_j^t| > |\mathcal{CP}_j^t| \\ \int_{-\infty}^{\infty} f_j^1(y) y dy & \text{if } |\mathcal{TP}_j^t| \leq |\mathcal{CP}_j^t| \end{cases}$$

where \bar{y} is the upper bound of the possible reserve price for resource I_j . The above estimation is "conservative" in the sense that we assume that agent a is less competitive than its trading competitors.

$\gamma(\mathcal{TAG}_j^t)$ models how the current set \mathcal{TAG}_j^t of agreements will affect agent a 's reserve price for resource I_j at round t . a will set a lower reserve price if it has made more agreements. Since current agreements may be decommitted in the future. Rather than considering the number $|\mathcal{TAG}_j^t|$ of agreements having already made, it's more prudent to use the expected number of final agreements, which can be computed based on the decommitment probabilities of agreement set \mathcal{TAG}_j^t . The decommitment probability of an agreement $Ag \in \mathcal{TAG}_j^t$ between a and s is approximated by considering the competition situation of negotiation over resource I_j and s 's satisfaction about the agreement Ag .

The competition situation of negotiation over resource I_j is evaluated by the probability that the agent s is not the most preferred trading partner is $[(\mathcal{TP}_j^t - 1)/\mathcal{TP}_j^t]^{c_{P_j^t}+1}$ [124, 125, 127]. s 's satisfaction about the agreement Ag is estimated by the probability that the agreement is no worse than the trading partner's reserve price. The price of the agreement $Ag \in \mathcal{TAG}_j^t$ is $\text{Prc}(Ag)$, s 's satisfaction about the agreement Ag is $F_j(\text{Prc}(Ag))$.

Hence, the approximation of the probability of s 's decommitting from agreement $Ag \in \mathcal{TAG}_j^t$ is defined as:

$$\omega_s^t(Ag) = \vartheta \times \left(1 - \left(\frac{\mathcal{TP}_j^t - 1}{\mathcal{TP}_j^t}\right)^{c_{P_j^t}+1}\right) (1 - F_j(\text{Prc}(Ag)))$$

For the tentative agreement set \mathcal{TAG}_j^t , the expected number of final agreements is $\varphi(\mathcal{TAG}_j^t) = \sum_{Ag \in \mathcal{TAG}_j^t} (1 - \omega_s^t(Ag))$. Given $\varphi(\mathcal{TAG}_j^t)$, buyer a can determine how it will affect the reserve price about resource I_j at round t . $\gamma(\mathcal{TAG}_j^t)$ decreases with the increase of $\varphi(\mathcal{TAG}_j^t)$ and can be defined as:

$$\gamma(\mathcal{TAG}_j^t) = \frac{1}{(1 + \varphi(\mathcal{TAG}_j^t))^2}$$

5.3.4 Maximum number of final agreements

Since trading partners may decommit from agreements, a may need to make more than one tentative agreement for resource I_j . Then, how many agreements are enough for the resource I_j ? For an agreement Ag between a and a trading partner s , s may be inclined to decommit if there are many buyers requesting the resource. On the other hand, s may be inclined to decommit if the agreement price is not favorable from s 's perspective. Here we provide an approach to decide the maximum number of agreements a can make on resource I_j at round t based on the expected number of final agreements. Given the expected number $\varphi(\mathcal{TAG}_j^t)$ of final agreements about resource I_j at t , a needs to decide whether the tentative agreements is enough or insufficient. If \mathcal{TAG}_j^t is more than needed, a may decommit from some agreements. If the agreement set is insufficient, a will make more agreements if the negotiation deadline hasn't approached. This work assumes that a only needs to make one final agreement for each resource. Therefore, by intuition, the most favorable result for agent a is that a makes exactly one final agreement for each resource.

As a only needs one final agreement about resource I_j , if $\varphi(\mathcal{TAG}_j^t) \gg 1$, only part of the final agreements will be used by a , which corresponds to the tentative agreement set $\mathcal{TAG} \subset \mathcal{TAG}_j^t$. Maintaining the tentative agreement set \mathcal{TAG} is better than maintaining the tentative agreement set \mathcal{TAG}_j^t as in the later case, a needs to pay more for redundant agreements. Therefore, it's better for a to decommit from some agreements in \mathcal{TAG}_j^t .

Let φ_j^t be the satisfactory number of final agreements about resource I_j at t which represents the upper bound of the number of final agreements needed. Before the deadline is reached, a has the opportunity to make more agreements and thus reach one final agreement. Thus, the satisfactory number of final agreements about resource I_j at $t < \tau$ is 1, $\varphi_j^t = 1$. After the negotiation deadline, a will determine whether to decommit from any agreement \mathcal{TAG}_j^t for resource I_j at round $\tau \leq t < t + \lambda$.

Is it the best option for a to set the satisfactory number of final agreements about resource I_j at t be 1? Consider the following scenario, at t , the expected number of final agreements for resource I_j is 1 and the expected number of final agreements about any other resource is close to 0, which implies that the negotiation about other resources has a very high failure probability. If a sets φ_j^t to be 1, it's with very high probability that a would need to decommit from all its agreements. Therefore, a will not set a high φ_j^t value if $\varphi(\mathcal{TAG}_k^t)$ is small for another resource I_k . On the other hand, a will try to increase the probability of making one final agreement for each resource as it's desirable for a to make one final agreement for each resource. Concerning above, φ_j^t is defined as:

$$\varphi_j^t = \begin{cases} 1 & \text{if } t < \tau \\ \min_{I_k \in \mathcal{I}^t} \varphi(\mathcal{TAG}_k^t) & \text{if } \tau \leq t < t + \lambda \end{cases}$$

If $\sum_{Ag \in \mathcal{TAG}_j^t} (1 - \omega_s^t(Ag)) < \varphi_j^t$, a needs to make more agreements as the expected number of agreements is less than φ_j^t . If $\sum_{Ag \in \mathcal{TAG}_j^t} (1 - \omega_s^t(Ag)) > \varphi_j^t$, a needs to decommit from some agreements. Let the set of tentative agreement set after removing unnecessary agreements be \mathcal{TAG} . The optimization problem of computing \mathcal{TAG} is given by

$$\min_{\mathcal{TAG}} \sum_{Ag \in \mathcal{TAG}_j^t - \mathcal{TAG}} \rho(\text{Prc}(Ag), \text{Tm}(Ag), t, \lambda)$$

where \mathcal{TAG} satisfies $\sum_{Ag \in \mathcal{TAG}} (1 - \omega_s^t(Ag)) \leq \varphi_j^t$.

Theorem 23. *The optimization problem of removing redundant tentative agreements is \mathcal{NP} -complete.*

Proof. We show that the problem is \mathcal{NP} -complete by formulating the problem as a 0-1 Knapsack problem, which is well known to be \mathcal{NP} -complete.

Formal definition of 0-1 Knapsack problem: There is a knapsack of capacity $c > 0$ and N items. Each item has value $v_i > 0$ and weight $w_i > 0$. Find the selection of items ($\delta_i = 1$ if selected, 0 if not) that fit, $\sum_{i=1}^N \delta_i w_i \leq c$, and the total value, $\sum_{i=1}^N \delta_i v_i$, is maximized.

The set of tentative agreements $\mathcal{TAG}_j^t = \{Ag_1, \dots, Ag_N\}$ can be treated as items. The value of each item Ag_i is defined as the penalty if a decommits from the agreement, i.e., $v_i = \rho(\text{Prc}(Ag_i), \text{Tm}(Ag_i), t, \lambda)$. The weight of each item Ag_i is defined as the probability that Ag_i will not be decommitted by a 's trading partner, i.e., $w_i = 1 - \omega_s^t(Ag_i)$. The capacity of the knapsack is defined as $c = \varphi_j^t$. $\delta_i = 1$ implies that Ag_i will be not decommitted by agent a .

The constraint of the optimization problem can be rewritten as the exact constraint $\sum_{i=1}^N \delta_i w_i \leq c$ of the 0-1 Knapsack problem. The optimization formula can be rewritten as

$$\min \sum_{i=1}^N (1 - \delta_i) v_i = \sum_{i=1}^N v_i + \min \sum_{i=1}^N -\delta_i v_i$$

which is equivalent to

$$\max \sum_{i=1}^N \delta_i v_i$$

Thus, the optimization problem can be formulated as a 0-1 Knapsack problem and it's \mathcal{NP} -complete. \square

A simple greedy approximation algorithm is used to compute the set of agreements which will not be decommitted by a (Algorithm 5) [44]: first sort all the tentative agreements \mathcal{TAG}_j^t by decreasing ratio of penalty to probability that an agreement will not be decommitted by a 's trading partners, then greedily pick agreements in this order (starting from the first agreement) until when adding a new agreement will violate the constraint of the maximum expected number of final agreements. For a removed agreement $Ag \in \mathcal{TAG}_j^t$, a decommits from the agreement; otherwise, a sends the agent s a proposal worse than $\phi_{s \rightarrow a}^t$.

Algorithm 5: Decommit from unnecessary agreements

Input: Tentative agreement set \mathcal{TAG}_j^t .

Output: Tentative agreement set \mathcal{TAG}_j^t satisfying the constraint of the maximum number of final agreements.

- 1: Sort all the tentative agreements \mathcal{TAG}_j^t by decreasing ratio of $\rho(\text{Pr}(Ag_i), \text{Tm}(Ag), t, \lambda)$ to $1 - \omega_s^t(Ag)$.
 - 2: Set $\mathcal{TAG} = \emptyset$, $i = 1$, and Ag be the i^{th} agreement in \mathcal{TAG}_j^t .
 - 3: **while** $\sum_{Ag' \in (\mathcal{TAG} + Ag)} (1 - \omega_s^t(Ag')) \leq \varphi_j^t$ **do**
 - 4: Add Ag into \mathcal{TAG} ;
 - 5: $i++$, and let Ag be the i^{th} agreement in \mathcal{TAG}_j^t ;
 - 6: **end while**
 - 7: **return** \mathcal{TAG}
-

5.4 Empirical evaluation and analysis

In this section, we first detail the methodology for analyzing the performance of the developed negotiation strategies. We then proceed to the actual empirical study of the proposed strategies. Finally, some properties of our negotiation strategies are analyzed.

5.4.1 The methodology

To evaluate the performance of negotiation agents, a simulation testbed consisting of a virtual e-Marketplace, a society of trading agents and a controller was implemented using JAVA. The controller generates agents, randomly determines their parameters (e.g., their roles as buyers or sellers, set of resources they provide or acquire, initial prices, reserve prices, deadlines), simulates the entrance of agents to the virtual e-Marketplace, and handles message passing and payment transfer.

5.4.1.1 Agent design

While there has been a lot of research in agent-mediated negotiation [70, 81, 90], most work focuses either on bilateral multi-issue negotiation (e.g., [49, 50, 51, 80, 81, 82, 133]) or single issue one-to-many negotiation (e.g., [13, 14, 28, 96, 97, 106]). One exception is [130] which studies concurrent one-to-many negotiations for multiple

resources. But in [130], an agent is assumed to know the reserve price of each resource. Given that there is no existing negotiation agents dealing with our multi-resource negotiation problem, for comparison reason, we implemented three other types of buyers based on existing techniques for single resource negotiation and negotiation with decommitment: 1) *TDAs* using a time-dependent strategy, 2) *MTDAs* using a market based time-dependent strategy, and 3) *ACMAs* using an adaptive commitment management strategy detailed in [130]. Experiments were carried out to study and compare the performance of our buyer agents (*HBAs*, heuristic-based buyer agents) with *TDAs*, *MTDAs*, and *ACMAs*.

TDAs, *MTDAs* and *ACMAs* adopt the strategy suggested by Nguyen and Jennings [98] and make at most one tentative agreement for each resource. *TDAs*, *MTDAs* and *ACMAs* use the same approach to determine the reserve price of each resource and use existing single resource negotiation strategies for the negotiation for each resource. The reserve price of resource I_j of each *TDA* (or *MTDA* and *ACMA*) is determined by considering the distribution of the reserve price of resource I_j . Specifically, the reserve price of resource I_j is proportional to its average reserve price. That is,

$$RP_j^t = RP^t \frac{\int_{-\infty}^{\infty} f_j(y)}{\sum_{i=1}^l \int_{-\infty}^{\infty} f_i(y)}$$

where l is the number of resources (i.e., issues) to acquire.

Similar to *HBAs*, *TDAs*, *MTDAs* and *ACMAs* generate proposals using a time-dependent negotiation decision function [48], which is widely used for designing negotiation agents (e.g., [13, 14, 48, 51, 97, 98, 124, 125, 127, 130]). However *TDAs*, *MTDAs* and *ACMAs* adopts different concession making strategies, i.e., they take different ε values. As *HBAs*, *TDAs* adopt the linear concession strategy, i.e., $\varepsilon = 1$. In contrast, *MTDAs* take market competition into account when making proposals. An *MTDA*'s parameter ε for concession making is adjusted in the following way: while the number of sellers are less than the number of buyers, an *MTDA* chooses

Table 5.2. Experimental settings

Input Data	Possible Values		
Market Type	<i>Favorable</i>	<i>Balanced</i>	<i>unfavorable</i>
<i>supply/demand</i>	10:1, 5:1, 2:1	1:1	1:2, 1:5, 1:10
Market Density	<i>Sparse</i>	<i>Moderate</i>	<i>Dense</i>
<i>No. of agents</i>	6 – 35	36 – 65	66 – 95
Deadline	<i>Short</i>	<i>Moderate</i>	<i>Long</i>
T_{max}	10 – 30	35 – 55	60 – 80
Resources/job	<i>Lower range</i>	<i>Mid-range</i>	<i>High range</i>
l	1 – 3	4 – 6	7 – 9

the conciliatory concession strategy by setting $\varepsilon < 1$. Otherwise, an *MTDA* uses the conservative or linear concession strategy by setting $\varepsilon \geq 1$. *MTDA*s' adaptive concession making strategy based on market competition has been shown to make *minimally sufficient* concessions in single resource negotiation [125]. *ACMA*s use the adaptive commitment management strategy used in [130] for each single resource negotiation. Specifically, *ACMA*s use a fuzzy decision making approach for deriving adaptive commitment management strategy profiles of buyers. The value of ε of a resource is determined dynamically at each round using fuzzy rules.

Each seller agent in the market randomly chooses a negotiation strategy from the set of alternations outlined in [48]: the time-dependent function (linear, conceder, conservative) and the behavior-dependent function (e.g., tit-for-tat). Each seller agent can only make at most one tentative agreement and it will decommit from an agreement if and only if it can benefit from the decommitment.

5.4.1.2 Experimental settings

In the experiments, agents were subjected to different market densities, market types, deadlines, number of resources to acquire or sell, and supply/demand ratio of each resource (see Table 5.2). Both market density and market type depend on the

probability of generating an agent in each round and the probability of the agent being a buyer (or a seller). When the number of agents are in the range of 6 – 35 (respectively, 36 – 65 and 66 – 95), the market is sparse (respectively, moderate and dense). The lifespan of an agent in the e-market, i.e., its deadline, is randomly selected from [10, 80]. The range of [10, 80] for deadline was adopted based on experimental tuning and agents’ behaviors. In our experimental setting, we found that: 1) for a very short deadline (< 10), very few agents could complete deals, and 2) for a deadlines longer than 80, there was little or no difference in the performance of agents. Hence, for the purpose of experimentation, a deadline between the range of 10 – 30 (respectively, 35 – 55 and 60 – 80) is considered as short (respectively, moderate and long). Each buyer may have different number of resources to acquire through negotiation. The number of resources each job (or task) needs is randomly selected from 1 to 9, where 1 – 3 (respectively, 4 – 6 and 7 – 9) is considered as lower range (respectively, mid-range and upper range). The value of ε (eagerness) is randomly generated from [0.1, 8] as it was found that when $\varepsilon > 8$ (respectively, $\varepsilon < 0.1$), there was little or no difference in performance of agents.

Each resource’s demand (i.e., the number of buyers who want to buy the resource) may not be equal to its supply (i.e., the number of sellers who want to sell the resource). If one buyer is negotiating for multiple resources, there are two situations: 1) All the resources have the same supply/demand ratio. From a buyer agent’s perspective, for a favorable (respectively, an unfavorable) market, the supply is much higher (respectively, lower) than the demand. 2) The resources have different supply/demand ratios. Then the range and variance of resources’ supply/demand ratios will affect agents’ performance. All our discussions of supply/demand ratio implicitly assume that the supply/demand ratio of each resource is randomly chosen.

There are four kinds of buyers (i.e., *HBA*, *TDA*, and *MTDA*, *ACMA*) and different kinds of sellers. The number of buyers (or sellers) of each kind is decided in a random

Table 5.3. Performance Measure

<i>Success Rate</i>	$R_{suc} = N_{success}/N_{total}$
<i>Expected Utility</i>	$U_{exp} = (\sum_{i=1}^{N_{total}} U_i)/N_{total}$
<i>Agreement per Resource</i>	$AG_{aver} = \frac{\sum_{i=1}^{N_{total}} \sum_{j=1}^{IS_i} A_i^j}{\sum_{i=1}^{N_{total}} IS_i}$
<i>Rate of Recovery from Decommittment</i>	$RR_{aver} = \frac{SD_{total}}{D_{total}}$
<i>Message per Resource</i>	$M_{aver} = \frac{\sum_{i=1}^{N_{total}} \sum_{j=1}^{IS_i} M_i^j}{\sum_{i=1}^{N_{total}} IS_i}$
N_{total}	<i>Total number of runs</i>
$N_{success}$	<i>No. of runs that reached consensus</i>
U_i	<i>Utility of the i^{th} run</i>
IS_i	<i>The number of resources in the i^{th} run</i>
A_i^j	<i>The number of tentative agreement for resource j in the i^{th} run</i>
M_i^j	<i>The number of messages for resource j in the i^{th} run</i>
D_{total}	<i>The number of runs in which one resource's tentative agreements were all decommitted</i>
SD_{total}	<i>The number of runs in which negotiation is successful after one resource's tentative agreements were all decommitted</i>

way. Without loss of generality, we assume that, there is at least one agent for each kind of agent.

5.4.1.3 Performance measure

We use a number of performance measures in the experiments (Table 5.3). Analyzing agents' utility can provide insights into how effective a strategy is. Since negotiation outcomes of each agent are uncertain (i.e., there are two possibilities: eventually reaching a consensus or not reaching a consensus), it seems more prudent to use expected utility for all runs (rather than expected utility for all successful runs) as a performance measure. For ease of analysis, agent a 's utility u_a (defined in Section 5.2.2) is normalized in each experiment in the following way: $u'_a = u_a/|RP_a - IP_a|$, which implies that $u'_a \leq 1$ if not considering the penalty a received from sellers. This

Table 5.4. Experimental results for 10^6 runs (performance measures are defined in Table 5.3)

<i>Strategy</i>	U_{exp}	R_{suc}	AG_{aver}	RR_{aver}	M_{aver}
<i>HBA</i>	0.206	0.59	1.34	$\frac{478}{1356} = 0.35$	86
<i>HBA-1</i>	0.153	0.58	1.27	$\frac{597}{2389} = 0.25$	91
<i>HBA-2</i>	0.111	0.50	1.33	$\frac{835}{3134} = 0.27$	89
<i>HBA-3</i>	0.144	0.43	0.63	$\frac{3052}{8945} = 0.34$	88
<i>HBA-12</i>	0.135	0.47	1.23	$\frac{2933}{9578} = 0.31$	92
<i>HBA-13</i>	0.144	0.42	0.63	$\frac{2704}{9362} = 0.29$	88
<i>HBA-23</i>	0.087	0.35	0.59	$\frac{3171}{10489} = 0.30$	84
<i>ACMA</i>	0.033	0.27	0.59	$\frac{3737}{13347} = 0.28$	85
<i>MTDA</i>	0.021	0.25	0.57	$\frac{4423}{15584} = 0.28$	84
<i>TDA</i>	0.019	0.25	0.72	$\frac{9200}{33459} = 0.27$	86

normalization is the same for agents with different strategies. It was pointed out in [70, 129] that in addition to optimizing agents' overall utility, enhancing the success rate is also an important evaluation criterion for designing negotiation agents.

In addition to the expected utility and success rate, it's necessary to compare the number of messages sent and received by each buyer during negotiation. As the number of resources each buyer is acquiring may be different at each time, it's intuitive to compare the number of messages sent or accepted for each resource. As an agent may make more than one tentative agreement for each resource, measuring the average number of tentative agreements for each resource is also important. During negotiation, it's possible that all of one agent's tentative agreements for one resource are decommitted by its trading partners and thus an agent's ability to recover from such situation is extremely important. Therefore, we also record and compare the number of cases where an agent makes a final agreement after all its tentative agreements for one resource are decommitted by its trading partners.

5.4.1.4 Results

A “matched-pair” study was conducted to evaluate the performance of *HBAs*, as compared with *TDA*s, *MTDA*s, and *ACMA*s. At the beginning of each run (experiment), the controller of the testbed will generate all the agents and set the parameters of all the agents according to the experimental setting, e.g., the number of agents, the supply/demand ratio of each resource, etc. Among all the buyers, there are some *target* buyers, one for each negotiation strategy we want to compare. All the target agents at each run have the same properties. For example, when we want to compare the performance of *HBAs* with *TDA*s, *MTDA*s, and *ACMA*s, we create one target *HBA*, one target *TDA*, one target *MTDA* and one target *ACMA*, which have the same properties (e.g., the set of resources to acquire, the reserve price, the initial price) except that they use different negotiation strategies. Then all the agents negotiate and compete with each other. At the end of this experiment, the controller will record the experimental results for each target agent, which will be averaged and analyzed on a large number of runs.

Extensive stochastic simulations were carried out for all the combinations of market density, market type and other agents’ characterizations. All the values of different performance measures were averaged based on more than 10^6 runs. In addition, we tried different decommitment deadlines and penalties functions. Even though experiments were carried out for all the situations, due to space limitations, only representative results are presented in this section. For the empirical results presented in this section, the market is of moderate density, $\lambda = 4$ is chosen as the decommitment period and the penalty function is $0.06 \times \text{Prc}(Ag) \times ((t' - t)/\lambda)^{1/2}$. $\lambda = 4$ is chosen based on the value of negotiation deadline. The shortest negotiation deadline is 10 in our experiments and setting a decommitment period shorter than negotiation deadline is reasonable. As in [5, 98], we choose a penalty function in which the penalty increases with the contract price and the period between agreement making

and decommitting. The multiplier 0.06 in the penalty function is chosen to make the decommitment penalty smaller than the contracting price. In the sensitivity analysis section (Section 5.4.2.6), we discussed the effect of changing the decommitment period and the penalty function. We also found that the confidence interval for each reported value is not wider than 0.001, which is negligible as compared with each value. Therefore, the experimental results can be trusted, which is mainly due to the large sample size.

5.4.2 Observations

5.4.2.1 Observation 1

HBA agents use three heuristics: Heuristic 1 (Section 5.3.2) is used to decide the deadline for each resource; Heuristic 2 (Section 5.3.3) is used to make a proposal for each resource in which the reserve price of each resource is dynamically chosen based on current market dynamics; Heuristic 3 (Section 5.3.4) is used to decide the number of tentative agreements to be made for each resource. Is it possible that a buyer in fact can get better negotiation performance by just using one or two heuristics? To verify that agents can get better negotiation performance by using all three heuristics simultaneously, we also compare the performance of *HBAs* with a special kind of buyers (called *HBA-s* here) which only use part of the heuristics used by *HBAs*. When a *HBA* doesn't use heuristic 2, it will use *MTDAs*' strategy to make proposals. When a *HBA* doesn't use heuristic 3, it makes at most one tentative agreement for each resource. *HBA-1s* are *HBAs* which don't use heuristic 1 and *HBA-12s* are *HBAs* which don't use heuristic 1 and heuristic 2. *HBA-123s* are equivalent to *MTDAs*. Table 5.4 shows the performance of *TDA*s, *MTDA*s, *ACMA*s, *HBA*s, and different types of *HBA-s* which only use part of *HBAs*' three heuristics.

From column 2 of Table 5.4, we can find that *HBAs* gain a higher expected utility U_{exp} than agents using other strategies. We also found that *HBA-s* get higher utilities

than *ACMAs*, *MTDAs*, and *TDAs*. In addition, heuristic 2 seems more important than the other two heuristics. *HBA-2s*' expected utility is lower than that of *HBA-1s* and *HBA-3s*. The average utility of *HBA-s* when *HBA-s* don't use heuristic 2 is $(0.111 + 0.135 + 0.087)/3 = 0.111$. The average utility of *HBA-s* when *HBA-s* don't use heuristic 1 is $(0.153 + 0.135 + 0.144)/3 = 0.144$. The average utility of *HBA-s* when *HBA-s* don't use heuristic 3 is $(0.144 + 0.144 + 0.087)/3 = 0.125$. Therefore, *HBA-s* will get lower utility when they don't use heuristic 2, as compared with not using either heuristic 1 or heuristic 3. In the same way, we can conclude that heuristic 3 is more important than heuristic 1. However, the above observations are based on the averaged results in *all* scenarios and they don't suggest that the heuristic 1 is more important than the other two heuristics in every specific scenario. When the supply/demand ratio of all the resources has a large variance, the average utility of *HBA-s* when *HBA-s* don't use heuristic 1 (respectively, heuristic 2 and heuristic 3) is 0.101 (respectively, 0.107 and 0.114), which implies that heuristic 1 is more important than the other two heuristics in this specific context. For all the values in columns 2 and 3, a t-test analysis with confidence level 95% was carried out and the difference between every two different values for the same performance is significant.


Column 3 of Table 5.4 shows that *HBAs* have higher success rates R_{suc} than agents using other strategies and *HBA-s* have higher success rates than *ACMAs*, *MTDAs*, and *TDAs*. In addition, heuristic 3 is more important than the other two heuristics from the perspective of achieving a higher success rate. This observation is intuitive since without using heuristic 3, each buyer makes only one tentative agreement and its probability of making a final agreement will be low if one or more trading partner decommits from an agreement. For the same reason, from column 4 of Table 5.4, we can see that *HBAs* have the highest number AG_{aver} of tentative agreements for each resource. *HBA-s* using heuristic 3 have more tentative agreements than *HBA-s* not using heuristic 3 which make at most one tentative agreement for each resource.

HBAs' number of runs in which all tentative agreements are decommitted is lower than all other kinds of buyers (see column 5 of Table 5.4). The recovery rate RR_{aver} of *HBAs* is also higher than the recovery rate of other kinds of buyers. For example, the recovery rate RR_{aver} of *HBAs* is $\frac{478}{1356} = 0.35$ indicating that there were 1356 situations in which all the tentative agreements for one resource were decommitted and 476 of the situations in which the agent made a final agreement. This observation corresponds with the intuition that *HBAs* are good at organizing and balancing the multi-resource negotiation. It's not surprising that *HBAs* will send more messages during negotiation as it may make more than one tentative agreement for each resource, which is mainly due to the use of heuristic 3. However, *HBAs*' average number M_{aver} of messages transferred for each resource is less than 3% higher than that of all other kinds of agents.

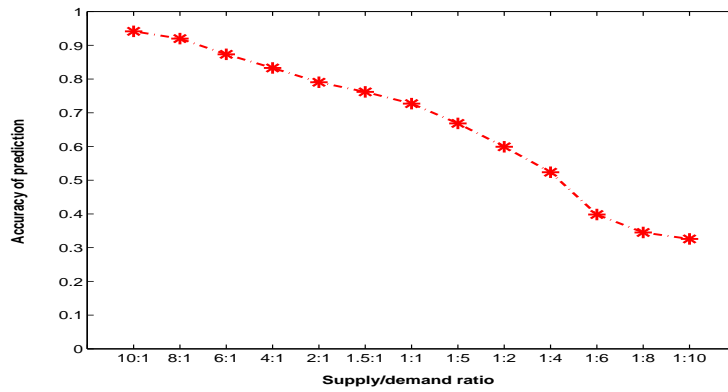
5.4.2.2 Observation 2

Our negotiation strategy uses the estimation of sellers' probability of decommitment. The decommitment probability is an approximation of the real probability, which is unknown to the buyer. It's impossible to justify our estimated "probabilities" with theory without making strong assumptions about knowing other agents' private information. Moreover, a seller's probability of decommitting from a tentative agreement is determined by many factors, e.g., its deadline, reserve price, its negotiation situation, which is unknown to a buyer.


Here we use an empirical approach to verify the accuracy of *HBAs*' estimation of decommitment probabilities. More specifically, *HBAs*' estimation of decommitment probabilities are compared with their trading partners' *real* decommitting actions during negotiation. Assume *HBAs* made n predictions $\langle \omega_s^t(Ag_1), \dots, \omega_s^t(Ag_n) \rangle$ for tentative agreements $\langle Ag_1, \dots, Ag_n \rangle$ throughout all the experiments in which $\omega_s^t(Ag_i)$ is a *HBA*'s predicted probability that its trading partner s will decommit from the agree-



(a) Prediction accuracy and deadline



(b) Prediction accuracy and market competition



(c) Prediction accuracy and the number of resources

Figure 5.2. Prediction accuracy of *HBAs*

ment Ag_i . Then *HBAs*' accuracy of predicting decommitting probabilities is given by:

$$AP = \frac{\sum_{1 \leq i \leq n} AP(\omega_s^t(Ag_i))}{n}$$

where

$$AP(\omega_s^t(Ag_i)) = \begin{cases} \omega_s^t(Ag_i) & \text{if } s \text{ decommits from } Ag_i \\ 1 - \omega_s^t(Ag_i) & \text{otherwise} \end{cases}$$

The average prediction accuracy in more than 10^6 runs is 0.774. Figure 5.2 shows the factors affecting the prediction accuracy. First, the prediction accuracy increases with the increase of *HBAs*' deadlines (Figure 5.2(a)). This result is intuitive as, with the increase of deadline, negotiation agents have longer time to interact with other agents. Then agents have a better understanding of the market and thus agents can make more precise predictions. Second, the prediction accuracy decreases with the increase of supply/demand ratio when all resources have the same supply/demand ratio (Figure 5.2(b)). When the supply/demand ratio is low, *HBAs* face high pressure of competition and decommitment is more likely to happen for each tentative agreement. As a consequence, it's more difficult to make a precise prediction. Finally, the prediction accuracy changes little with the change of the number of resources (Figure 5.2(c)). This observation is also intuitive as, a seller's decommitment decision is only affected by the agreement price, its reserve price and market competition. It has nothing to do with the negotiation status of other resources.

Our function of decommitment probability is based on our intuitions about which factors affect agents' decision to decommitment. The parameter $\vartheta = 0.68$ is a parameter of the function for computing trading partners' decommitment probabilities, which is based on experimental tuning. With the experimental tuning, we were able to get 77.4% accuracy averaged over all environments. However, it is unclear to us whether we can get a better result considering that *HBAs* do not know other agents' strategies nor their exact reserve prices. On a more positive note, our heuris-


tic function performs in ways that would be expected. For instance, when a *HBA*'s uncertainty reduces (e.g., change the distribution of sellers' reserve prices), it gains higher prediction accuracy. A reasonable prediction approach should have the property that the prediction accuracy increases with the decrease of uncertainty which our approach does. Although we can reduce uncertainty in the market and thus get higher prediction accuracy, our experiments will become less interesting. In addition, it's impractical to assume that agents have (almost) complete information about others.

Sim *et al.* [128, 130] also proposed a function for evaluating a trading partner's decommitting probability, which are used by *ACMAs* and achieved an average 38% accuracy in all the scenarios. Although the function in [130] appears to be simpler as it only considers the prices of the proposals it has received, it is noted that [130] did not make the assumption that an agent has knowledge of the number of competitors. In contrast, our function takes both market competition and the trading partner's satisfaction of agreements based on each agent's knowledge about 1) the number of trading competitors and 2) the reserve price of each trading partner.


5.4.2.3 Observation 3

The experimental results in Figure 5.3 show that: 1) Negotiation results become more favorable with the increase of the deadline for all kinds of buyers. With short (respectively, long) deadlines, different kinds of agents have equally insufficient (respectively, sufficient) time to optimize their agreements. 2) Given the same deadline, *HBAs* achieved higher utilities than *ACMAs*, *MTDAs*, and *TDA*s. 3) The advantages of *HBAs* over *MTDAs* and *TDA*s decreases when the market becomes more favorable.


Experimental results in Figure 5.4 indicate that the success rate of *HBAs* are always higher than that of *ACMAs*, *MTDAs*, and *TDA*s. However, this advantage decreases when the market become more favorable. In addition, with the increase of deadline, agents' success rates have a large increase at the beginning and slightly



(a) Unfavorable market




(b) Balanced market




(c) Favorable market


Figure 5.3. Deadline and expected utility



(a) Unfavorable market




(b) Balanced market




(c) Favorable market


Figure 5.4. Deadline and success rate



(a) Short deadline




(b) Moderate deadline




(c) Long deadline


Figure 5.5. Number of resources and expected utility



(a) Short deadline




(b) Moderate deadline




(c) Long deadline


Figure 5.6. Number of resources and success rate



(a) Short deadline




(b) Moderate deadline




(c) Long deadline


Figure 5.7. Supply/demand ratio and expected utility



(a) Short deadline



(b) Moderate deadline



(c) Long deadline

Figure 5.8. Supply/demand ratio and success rate

decrease when the deadlines are long. When agents have long deadlines, agents have more time to bargain with other agents and seek good agreements with the increase of deadlines. Since agents use time-dependent strategies, buyers with longer deadlines are inclined to make less concessions at each time as agents will prefer to propose their reserve prices when their deadlines approach. Thus, buyers will become more patient and will not accept proposals which are not favorable enough while considering their future opportunities to make better agreements. Therefore, buyers with longer deadlines will fail to make agreements with some sellers, especially sellers with shorter deadlines. Although buyers' success rates decrease with the increase of deadlines when deadlines are relatively long, buyers' utilities increase with the increase of deadlines. This is because buyers will set higher expectation about the agreements with the increase of deadlines. Thus, the agreements made by buyers with longer deadlines are more favorable as compared with agreements made by buyers with shorter deadlines.

5.4.2.4 Observation 4

From Figure 5.5 we can see that, as the number of resources to be acquired increases, the utilities of all kinds of agents decrease. That is because, with the increase of the number of resources each agent acquires, it's harder to manage all the negotiations and the probability that the overall negotiation fails increases, which directly correlates with the decreased success rates in the strategies explored here. *HBAs* always achieved higher utilities than *ACMAs*, *MTDAs*, and *TDAs*.

Experimental results in Figure 5.6 indicate that the success rate of *HBAs* are always higher than that of *ACMAs*, *MTDAs*, and *TDAs*. However, this advantage decreases when agents have longer deadlines as in this case, all agents have enough time to negotiate for agreements. Agents' success rate decreases significantly as a small number of resources (e.g., 1 or 2). With more resources, it's more difficult for

buyers to manage and establish agreements for all resources because of the difficulties of managing all the negotiation threads.

5.4.2.5 Observation 5

It can be observed from Figure 5.7 and Figure 5.8 that *HBAs* always get higher utilities (respectively, success rates) than *ACMAs*, *MTDAs*, and *TDA*s when all resources have the same supply/demand ratios. Additionally, when the supply/demand ratio is high (e.g., 10), the average utilities of the three types of agents are close especially in the long deadline case since agents have many choices and can easily switch from one agreement to another agreement, i.e., there is limited space to optimize the agreements. The advantage of *HBAs* in success rate decreases when agents have longer deadlines. Since buyer agents with different strategies compete with each other, it is possible that one strategy achieved much better negotiation results than another strategy in a specific market. Due the strategic interaction among agents, one strategy may achieve a good performance in only certain markets. In Figure 7 we can see that when the ratio is in the range 0.5-0.7, *MTDAs* achieved very low utilities as compared with the utilities when the ratio less than 0.5 or higher than 0.7. When the ratio is in the range 0.5-0.7, *HBAs* achieved higher utility than that when the ratio less than 0.5 or higher than 0.7. When the supply/demand ratio is very low (e.g., 0.2-0.4), it is difficult for an agent to get agreements, thus all different strategies achieved low utilities. When the supply/demand ratio are slightly low (e.g., 0.5-0.7), some *HBAs* may make agreements for all required resources. An *MTDA* can also make agreements for some of its resources using its market-driven concession strategy. However, since *MTDAs* are lacking of the ability of coordinating their negotiation for multiple resources. They often can only satisfy part of their resources. Therefore, when the whole negotiation failes, an *MTDA* either pays a lot of penalties to decommit from its agreements or pays for some final agreements which have not

been decommitted. Accordingly, *MTDAs* often get negative utilities. When *MTDAs* decommit from agreements, *HBAs* have a better chance to make new agreements in this situation. The experimental results also show that when the supply/demand is in the range 0.5-0.7, *MTDAs* made more agreements (including both tentative and final) than *TDA*s and *ACMA*s but the success rate of *MTDAs* is not higher than that of *TDA*s and *ACMA*s. When the market is almost balanced (e.g., the supply/demand is in the range of 0.8-1), it is easier for *MTDAs* to make agreements which can satisfy their resource requirements and their utilities are much higher than that when the supply/demand is in the range 0.5-0.7.

5.4.2.6 Sensitivity analysis

We also did additional experiments to explore how sensitive are our experimental results to changes of the parameters of our experimental environments or assumptions about our negotiation model.

1) With the increase of penalty, the average utility of agents including *HBAs* decreases. For example, when we double the penalty fee, the average utility of *HBAs* is decreased by 7%. The main reason is that with a higher penalty, a buyer is more likely to commit to an early agreement, which may have a low utility value. When the penalty fee is low, a buyer will decommit from an early agreement and make a new agreement with a higher utility value. Similarly, each seller is also more likely to stick to an early agreement when the penalty is high. *HBAs* always have better performance than other types of buyers when using different penalty functions.

2) With the increase of decommitment period λ , the average utility of agents including *HBAs* decreases. For instance, when we set a decommitment period $\lambda = 6$ instead of 4, the average utility of *HBAs* is decreased by 8%. With a longer decommitment period, the probability that an agreement will be decommitted will increase, and thus the probability that a buyer will get a final agreement decreases. However,

the advantage of *HBAs* over other types of buyers increases with the increase of the decommitment period λ as buyers like *ACMAs*, *MTDAs*, and *TDA*s make at most one tentative agreement for each resource.

3) When agents have more accurate information about other agents, agents including *HBAs* achieved better performance. This chapter assume that a buyer knows the probability distribution of sellers' reserve prices and the number of competitors. We find that that the accuracy of this information does have an effect on agents' negotiation performance. When a buyer's knowledge becomes less accurate, its utility decreases. For example, when the believed number of competitors is less than half of the actual number of competitors, the average utility of *HBAs* is 7% lower than that of *HBAs* knowing the actual number of competitors. However, even with this level incorrect information, *HBAs* still achieved better performance than other types of agents.

4) While keeping the supply/demand ratio of each resource constant, market density has little effect on agents' performance. In a moderate density market, agents' average utilities are 2% lower than that in a market of dense density and are 1% higher than that in a market of sparse density.

5.4.3 Analysis of properties

Typically, agents use a *monotonic* concession protocol by insisting on their previous proposals or raising/reducing their proposals monotonically until an agreement is reached. In a dynamic negotiation environment, market competition and agents' evaluation may change over time, protocols that are not monotonic may achieve higher average utilities. Negotiation agents in this chapter make a proposal based on market situation and the negotiation situations of other threads. Therefore, the proposed negotiation protocol is not monotonic.

In a favorable market, there are fewer competitors and more trading partners. Hence, an agent has stronger bargaining power and doesn't need to make large concessions. In an unfavorable market, an agent experiences more competition, and it may attempt to make more concessions. With respect to competition, an agent strives to avoid making large concessions in favorable markets or making too large concessions in unfavorable markets. Additionally, when the expected number of final agreements is high, an agent is inclined to make less concession as it only needs one final agreement.

Property 24. *Agents will make less concession with the increase of the expected number of final agreements when the worst possible utility doesn't increase.*

Take the resource I_j for example. The number of agreements has no effect on the expected agreement price ϖ_j^t . As the worst possible utility doesn't increase, the conflict probability χ_j^t will not increase. $\gamma(\mathcal{TAG}_j^t)$ will decrease with the expected number of final agreements $\varphi(\mathcal{TAG}_j^t)$. Therefore, the reserve price of resource I_j will decrease and thus agents will make less concession.

Property 25. *Agents will make less (respectively, more) concession with the increase of the number of trading partners (respectively, competitors).*

Take resource I_j for example. The number of trading partners has no effect on $\gamma(\mathcal{TAG}_j^t)$. With the increase of trading partners, χ_j^t will not increase and ϖ_j^t will also not increase. Thus, the reserve price of resource I_j will not increase and thus agents make less concessions. Similarly, with the increase of trading competitors, ϖ_j^t will decrease. Thus, the reserve price of resource I_j will increase and thus agents will make more concession.

Property 26. *When competition is high and penalty is very low, agents may make agreements with all the trading partners.*

Take resource I_j for example. The decommitment probability increases with the increase of competition. As the penalty is very low, an agent with more tentative agreements won't pay too much penalty when it has to decommit from some tentative agreements. An extreme situation is that the agent can even make agreements with all the trading partners.

From Properties 2 and 3 we can learn that the market competition places an important role on deciding the amount of concessions and the number of tentative agreements. With respect to competition, a negotiation agent decides the maximum number of agreements. In a favorable market, there are fewer competitors and more trading partners. Hence, an agent doesn't need to make many agreements (concessions, respectively). In an unfavorable market, an agent's bargaining power decreases as it experiences more competition, and it may attempt to make more agreements (concessions, respectively) as its trading partners are more likely to decommit from agreements.

One possible strategy is to make agreements later and thus potentially a buyer will pay less decommitment penalties given that the penalty will increase with time. However, "delaying" agreements will also increase the probability that the whole negotiation fails. In addition, generally a buyer will increase its offering price gradually and it is possible that it can get some resources with a cheap price in the early negotiation stages. While taking the "delaying" strategy, the buyer will miss those cheap resources and buy expensive resources in a later time. Another disadvantage of delaying agreements is that the buyer may fail to get all resources when one seller decommits from agreements when the deadline is approaching. In our model, no agent can decommit from an agreement after a fixed time period based on when the agreement was made. Accordingly, making agreements earlier can potentially avoid negotiation's "collapsing" at the last minute. We examined agents' performance when

they choose delaying agreements and found that such strategic “delaying” do not improve agents’ performance.

As a result of this extensive empirical analysis, we have verified that the negotiation strategy for multi-resource acquisitions is both very effective in comparison to existing approaches and behaves in a consistent and appropriate manner as important characteristics of the marketplace are varied.

5.5 Summary

This chapter presents the design and implementation of negotiation agents that negotiate for multiple resources where agents don’t know the reserve price of each resource and are allowed to decommit from existing agreements. The contributions of this chapter include: 1) To avoid the risk of the “collapse” of the overall negotiation due to failing to acquire some scarce resources, negotiation agents have the flexibility to adjust the deadline for different resources based on market competition, which allows agents to response to uncertainties in resource planning. 2) Each agent utilizes a time-dependent strategy in which the reserve price of each resource is dynamically determined by considering (*conflict probability*), *expected agreement price*, and *expected number of final agreements*. 3) As agents are permitted to decommit from agreements, an agent can make more than one agreement for each resource and the maximum number of agreements is constrained by the market situation. 4) An extensive set of experiments were carried out and the experiments results show that each of the proposed heuristics contributes to improve agents’ performance and our proposed approach achieved better negotiation results than representative samples of existing negotiation strategies.

The experimental results showed that *HBAs* achieved better negotiation results (higher expected utilities and higher success rates) than *ACMAs*, *MTDAs*, and *TDA*s. Moreover, it’s better for *HBAs* to use all the three heuristics together as each heuris-

tic has different features. The heuristic for proposal creation seems more important than the other two heuristics. From our experimental results we can see that, when the negotiation environment is either very “tough” (i.e., short deadline, high competition, and more resource to negotiate) or very “favorable” (i.e., long deadline, less competition, and less resource to negotiate), *HBAs* did not significantly outperform *MTDAs* and *TDA*s. That is because in a “tough” market, all the agents have little opportunity for making individual agreements, and thus it’s very hard to find a good set of agreements that satisfy all the resource requirements. In contrast, in a very “favorable” market, agents can easily make good agreement set. It is in the middle ground that you see the significant advantage of our approach.

CHAPTER 6

NEGOTIATION WITH DECOMMITMENT FOR DYNAMIC RESOURCE ALLOCATION IN CLOUD COMPUTING

We consider the problem of allocating networked resources in dynamic environment, such as cloud computing platforms, where providers strategically price resources to maximize their utility. Resource allocation in these environments, where both providers and consumers are selfish agents, presents numerous challenges since the number of consumers and their resource demand is highly dynamic. While numerous auction-based approaches have been proposed in the literature, this work explores an alternative approach where providers and consumers automatically negotiate resource leasing contracts. Since resource demand and supply can be dynamic and uncertain, we propose a distributed negotiation mechanism where agents negotiate over both a contract price and a decommitment penalty, which allows agents to decommit from contracts at a cost. We compare our approach experimentally, using representative scenarios and workloads, to both combinatorial auctions and the fixed-price model used by Amazon’s Elastic Compute Cloud, and show that the negotiation model achieves a higher social welfare. Different from designing negotiation strategies to maximize an agent’s utility in the previous chapter, the focus in this chapter is designing a mechanism to maximize the social welfare (i.e., the sum of all agents’ utilities). The cloud computing resource allocation problem in this chapter is different from the resource allocation problem in the previous chapter in a number of aspects. For instance, a seller can sell multiple resources to a buyer and thus the buyer needs to consider different “plans” (each plan specifies the set of resources to

buy from different sellers) to satisfy its resource requirement. In addition, sellers need to provide resources during a fixed time period. These differences lead to the need for new approaches that were not considered in the Chapter 5. Before introducing the cloud computing problem and our resource allocation mechanism, we first discuss the idea of negotiating over penalty and its advantages.

6.1 Negotiation Over Decommithment Penalty

In automated negotiation systems for self-interested agents, contracts have traditionally been binding and do not allow agents to efficiently deal with future events in the environment. Sandholm and Lesser [121] proposed leveled-commitment contracts which allow an agent to be freed from an existing contract at the cost of simply paying a penalty to the other contract party. A self-interested agent will be reluctant to decommit because the other contract party might decommit, in which case the former agent gets freed from the contract, does not incur a penalty, and collects a penalty from the other party. Despite such strategic decommitting, leveled-commitment increases the expected payoffs of all contract parties and can enable deals that are impossible under full commitment [121]. This approach has been applied in a number of different applications [5, 11, 97, 98].

In leveled-commitment contracting, both contract parties strategically choose their level of commitment based on the contract price and decommitment penalty which are determined prior to the start of the decommitting game. The efficiency of leveled-commitment contracting depends on how the contract price and decommitment penalty are set. In Sandholm *et al.*'s model of leveled-commitment contracts [116, 121, 122], both the contract prices and decommitment penalties are assumed to be known to the contract parties before the decommitting game. Although algorithms are provided to optimize the social welfare of the equilibrium outcome [116], the optimization is not for the favor of each contract party. In existing applications

(e.g., [5, 11, 97, 98]) of automated negotiation with decommitment, decommitment penalties are set by third parties (e.g., system designers) and are either fixed or a function of contract prices.

Negotiating simultaneously over contract prices and decommitment penalties is appropriate for several reasons. First, it is difficult for system designers to decide *optimal* contract prices and decommitment penalties to maximize the social welfare, especially when there are multiple agents and agents have incomplete information. It is also intractable to compute agents' rational equilibrium strategies in many practical sequential games. Furthermore, it is not appropriate to assume that system designers have complete knowledge about agents in the system. Finally, a selfish agent may feel it is advantageous for it to decide the contract price and penalty by itself. When agents are allowed to negotiate over penalties, each agent has a larger strategy space which gives it more options for how to react to the current situation and it may be able to achieve a utility which cannot be achieved when it is not allowed to negotiate over penalty.

This section analyzes agents' strategic behavior in the bilateral contracting game prior to the decommitting game to make agreements on a contract and a decommitting penalty. One selfish contract party may prefer another pair of contract price and decommitting penalty to the contract price and decommitment penalty which maximize the social welfare. The leveled-commitment contracting we propose includes two games: a contracting game where the two parties bargain over contract price and decommitment penalty and a decommitting game in which the two agents make strategically decommitting decisions. During the decommitting game, agents will make optimal decommitting decisions while taking into account the contract price and decommitment penalty previously agreed upon. Therefore, in the contracting game, each agent will try to make the best contract price and penalty which will maximize its utility in the decommitting game.

As in [116, 121, 122], we consider a contracting setting with two risk neutral agents who attempt to maximize their own expected payoff: contractor \mathbf{b} who pays to get a task done, and contractee \mathbf{s} who gets paid for handling the task. The setting can be interpreted as modeling a variety of scenarios, for example bargaining between a buyer and a seller in e-commerce. In our model, \mathbf{b} and \mathbf{s} negotiate over contract price and decommitment penalty before additional offers (outside offers) from other agents become available. Then they strategically choose to decommit or not when their outside offers are available.

6.1.1 Leveled-commitment contracting

We study a setting where the future of agents involves uncertainty. We model this as agents' potentially receiving outside offers as in [116, 121, 122]. The contractor's outside offers could come from some other contractees which can provide the service requested by the contractor. The contractor can make agreements with those contractees in the future. The contractor's best (lowest) outside offer v is characterized by a probability density function $f(v)$. The contractee's best (highest) outside offer w is characterized by a probability density function $g(w)$. $f(v)$ and $g(w)$ are assumed statistically independent and are common knowledge [116, 121, 122]. That is, both agents have symmetric information as they both don't know the value of v and w .

The contractor's options are either to make a contract with the contractee or to wait for future option v . Similarly, the contractee's options are either to make a contract with the contractor or to wait for future option w . The two agents could make a full commitment contract at some price. Alternatively, they can make a leveled-commitment contract which is specified by a contract price, ρ , and a decommitment penalty q . If one agent decommits from the agreement, it needs to pay the

penalty q to the other agent.¹ When the decommitment penalty q is very large, a leveled-commitment contract is equivalent to a full contract as no agent will choose to decommit. Therefore, full commitment contracts are a subset of leveled-commitment contracts.

One implicit assumption is that during the contracting game, the contractor can only bargain with one contractee and the contractee can also only negotiate with one contractor. The other assumption is that the bargaining game finishes before outside options become available. Bargaining protocols can be used to control the length of negotiation. Moreover, even if agents are allowed to conduct infinite time negotiation, negotiation often stops soon since bargaining agents usually have deadline constraints and often face bargaining costs.

The leveled-commitment contracting consists of two stages. In the first stage, which we call the *contracting game*, the agents make agreements on both a contract price and a decommitting penalty. In the second stage, which we call the *decommitting game*, the agents decide on whether to decommit or not. Clearly, the equilibrium of the decommitting game affects the agents' preferences over contract prices and decommitment penalties in the contracting game. There is no decommitting game if agents make a **null** contract (i.e., no agreement is made) in the contracting game.

6.1.1.1 Contracting game

We consider the widely used one-shot protocol [111]. Formally, agent $\mathbf{a} \in \{\mathbf{b}, \mathbf{s}\}$ makes an offer $[\rho, q]$ where ρ is contract price and q is decommitment penalty. The other agent $\hat{\mathbf{a}}$ can choose to 1) **accept** or 2) **reject**. If $\hat{\mathbf{a}}$ accepts the offer, the

¹Our analysis can be easily extended to handle the setting where the penalties for the contractor and the contractee are different. Setting different penalties for contractor and contractee only makes it difficult to solve the decommitment game, which has been thoroughly analyzed in the work by Sandholm *et al.* [121]. For the contracting game, a new variable will be added but the analysis will be the same.

bargaining outcome is $[\rho, q]$. Otherwise, the bargaining fails and the outcome is **null** contract.

6.1.1.2 Decommitting game

The decommitting game happens only when the two agents make a leveled-commitment contract $[\rho, q]$. In the decommitting game, each agent has exactly one chance to decommit and there are different decommitting mechanisms depending on who decommits first [116, 121, 122]: 1) contractee has to reveal its decision first; 2) contractor has to reveal its decision first; and 3) agents reveal their decisions simultaneously. We consider the first decommitting mechanism, i.e., contractee takes actions first and contractor moves next. The other mechanisms can be analyzed analogously.

6.1.2 Optimal contracts

Agents' bargaining strategies in the contracting game are affected by the outcome of the decommitting game: each agent wants to make the optimal contract that maximizes its expected utility in the decommitting game. There may be multiple optimal contracts or the optimal contract may be the **null** contract.

We follow the same analysis as in [121] to compute agents' optimal contracts. Assume that the contract made during the contracting game is $[\rho, q]$. In a sequential decommitting game where the contractee has to decommit first, if the contractee has decommitted, the contractor's best move is not to decommit as $q \geq 0$. In the subgame where the contractee has not decommitted, the contractor's best move is to decommit if $-v - q > -\rho$, i.e., the contractor decommits if its outside offer, v , is below a threshold $v^* = \rho - q$. So, the probability that it decommits is $p_{\mathbf{b}} = \int_{-\infty}^{v^*} f(v)dv$.

The contractee gets $w - q$ if it decommits, $w + q$ if it does not but the contractor does, and ρ if neither decommits. Thus the contractee decommits if $w - q > p_{\mathbf{b}}(w + q) + (1 - p_{\mathbf{b}})\rho$. When $p_{\mathbf{b}} < 1$ the inequality above shows that the contractee decommits

if its outside offer exceeds a threshold $w^* = \rho + q(1 + p_{\mathbf{b}})/(1 - p_{\mathbf{b}})$. So, the probability that it decommits is $p_{\mathbf{s}} = \int_{w^*}^{\infty} g(w)dw$.

Given agents' equilibrium strategies under contract $c = [\rho, q]$, \mathbf{b} 's expected payoff $\pi_{\mathbf{b}}(c, f(v), g(w))$ ($\pi_{\mathbf{b}}(c, f, g)$ for short) is

$$p_{\mathbf{s}} \int_{-\infty}^{\infty} (q-v)f(v)dv + (1 - p_{\mathbf{s}}) \left[\int_{-\infty}^{v^*} -(v+q)f(v)dv - \int_{v^*}^{\infty} \rho f(v)dv \right]$$

The expected payoff $\pi_{\mathbf{s}}(c, f, g)$ of contractee \mathbf{s} is

$$\int_{w^*}^{\infty} g(w)(w - q)dw + \int_{-\infty}^{w^*} g(w) [p_{\mathbf{b}}(w + q) + (1 - p_{\mathbf{b}})\rho] dw$$

If agents fail to make a contract, an agent can wait for its best outside offer. Thus, agents' expected utilities under the **null** contract are $\pi_{\mathbf{b}}(\mathbf{null}, f, g) = \int_{-\infty}^{\infty} -f(v)v dv = -E(v)$ and $\pi_{\mathbf{s}}(\mathbf{null}, f, g) = \int_{-\infty}^{\infty} g(w)w dw = E(w)$.

Based on this analysis developed previously by Sandholm *et al.* [116, 121, 122], we now extend it to a contracting game and discuss agents' optimal contracts when agents negotiate over contract prices and decommitment penalties. We assume that agents are *individually rational* (IR), i.e., no agent will accept a contract worse than the **null** contract. A contract c is IR if it is individually rational for both agents. Formally, the set $\mathcal{C}(f, g)$ of IR contracts based on agents' beliefs $f(v)$ and $g(w)$ are

$$\{c | \pi_{\mathbf{b}}(c, f, g) \geq -E(v), \pi_{\mathbf{s}}(c, f, g) \geq E(w)\}$$

We assume that $\mathcal{C}(f, g)$ is not empty. The contract $c_{\mathbf{b}}^*(f, g)$ ($c_{\mathbf{s}}^*(f, g)$) which maximizes the contractor's (contractee's) expected utility is the contractor's (contractee's) optimal contract. Formally,

$$c_{\mathbf{b}}^*(f, g) = \operatorname{argmax}_{c \in \mathcal{C}(f, g)} \pi_{\mathbf{b}}(c, f, g)$$

$$c_s^*(f, g) = \operatorname{argmax}_{c \in \mathcal{C}(f, g)} \pi_s(c, f, g)$$

Therefore, the utility \mathbf{a} can get is in the range $[\pi_{\mathbf{a}}(\mathbf{null}, f, g), \pi_{\mathbf{a}}(c_{\mathbf{a}}^*(f, g), f, g)]$.

6.1.3 Efficiency of Negotiating Over Penalty in Two-player Game

In this section, we experimentally evaluate the efficiency of negotiating over penalty in the two-player game considered by Sandholm *et al.* [116, 121, 122]. Each game is characterized by contractor's best (lowest) outside offer v is characterized by probability density functions $f(v)$ and $g(w)$. For each game, we compare the contracting results when agents' decommitment penalties are determined by negotiation and results where decommitment penalties are determined exogenously. While there are many methods to exogenously set decommitment penalties [5, 11, 15, 97, 98], the following two approaches are the most widely used: 1) fixed penalty independent of contract prices and 2) penalty as a percentage of contract prices. We compare our negotiation based approach with the above two approaches. For fixed penalty, the penalty is chosen from $\{0, 10, 20, 40\}$. When the decommitment penalty is a percentage of a contract price, the rate is chosen from $\{0.1, 0.3, 0.5\}$, i.e., $q = 0.1\rho$, $q = 0.3\rho$ or $q = 0.5\rho$. Thus, there are 8 approaches to set penalties: 4 fixed penalty values, 4 penalty functions in which a penalty is a fraction of contract price, and the bargaining approach. For the bargaining approach, one agent is randomly chosen to offer to the other agent.

When the penalty is set exogenously, it is possible that two agents fail to make an agreement whereas they can make an agreement while negotiating over penalty. For instance, when both v and w are uniformly distributed between 30 and 90, the two agents can make an agreement when they are negotiating over penalty. However, they cannot make any agreement when the penalty is $q = 40$ or $q = 0.5\rho$ for the special distributions, no matter which agent is the offering agent.

We extensively evaluated the average social welfare (i.e., the sum of both agents' average utilities) of different mechanisms in a variety of settings. We found that the negotiating over penalty achieved higher social welfare than other penalty setting approaches. Figure 6.1 shows the performance of different mechanisms as well as the maximum social welfare (which is achieved when the offering agent or a third party chooses to maximize the social welfare rather than its utility) when $f(v)$ and $g(w)$ are uniform distributions. $f(v)$ is defined by $[v_{min}, v_{max}]$ and $g(w)$ is defined by $[w_{min}, w_{max}]$ where 1) $0 < v_{min}, v_{max}, w_{min}, w_{max} \leq 100$ and 2) $v_{max} \geq w_{min}$. That is, the outside offers v and w are in the range of $(0, 100]$ with the constraint that the contractor's outside offer v is not always less than the contractee's outside offer w .² All values reported in Figure 6.1 is the average value in 10000 randomly generated settings. We can see that negotiating over penalty achieved much higher utility than other exogenous penalty setting mechanisms. It can also be seen from Figure 6.1 that, even when the offering agent always chooses the price and penalty to maximize its utility, the social welfare is close to the maximum social welfare.

Figure 6.2 shows the performance of different mechanisms under the same setting except $50 < v_{min}, v_{max} \leq 100, 0 < w_{min}, w_{max} \leq 50$. That is, the contractor's outside offer v is always no less than the contractee's outside offer w . In this situation, the optimal offer of the offering agent is the pair so that no agent has an incentive to deviate from the contract. Consider that the contractor is the offering agent as an example. The optimal price the contractor is $E(w)$ and its optimal penalty q is one value such that the contractee has no incentive to decommit. Obviously, the contractor will also not decommit since $E(w) < E(v)$. If the contractor offers a price lower than $E(w)$, the contractee will not accept it. The corresponding social welfare is then 0. In this situation, the maximum social welfare which can be achieved by

²Note that if $v < w$, the two agents cannot make a contract.




Figure 6.1. Efficiency comparison in two-player game (1)

exogenously setting the offer is always 0 when they make an agreement. We can find that negotiating over penalty can always achieve the maximum social welfare. We also noticed from Figure 6.2 that by exogenously setting a high penalty (40 in this special case), agents can also achieve the maximum social welfare.

We have shown that in the canonical two player leveled decommitment games, negotiating over penalty achieved higher social welfare than exogenous penalty setting mechanisms. In the two player games, agents have symmetric information and the offering agent is able to compute its optimal offer by solving the contracting games and decommitment games. In more realistic scenarios in which there are usually more than two agents and agents have more uncertainties (e.g., outside options), it may be intractable to compute the optimal penalty to optimize the social welfare. Strong assumptions in the two player games leads us to investigate the benefits of negotiating over penalty in more realistic scenarios.

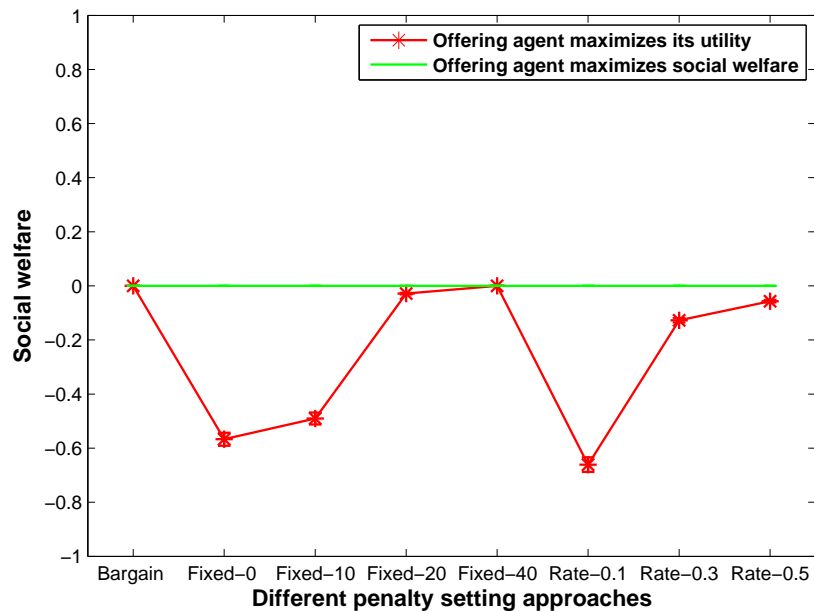


Figure 6.2. Efficiency comparison in two-player game (2)

We have shown that in the canonical two player leveled decommitment games, negotiating over penalty achieved higher social welfare than exogenous penalty setting mechanisms. In the two player games, agents have symmetric information and the offering agent is able to compute its optimal offer by solving the contracting games and decommitment games. In more realistic scenarios in which there are usually more than two agents and agents have more uncertainties (e.g., outside options), it may be intractable to compute the optimal penalty to optimize the social welfare. Strong assumptions in the two player games leads us to investigate the benefits of negotiating over penalty in more realistic scenarios such as the resource allocation problem in GENI.

6.2 Resource Allocation in GENI

We explore our resource allocation problem in the context of NSF's GENI initiative [1], which is building a prototype of a shared experimental infrastructure to

investigate next-generation Internet applications. GENI is similar to other cloud platforms in that it exposes network-accessible APIs for consumers to lease virtualized hardware components, although GENI offers a more diverse collection of resources donated by many providers, such as universities and industry research labs. The intent is for researchers to experiment with new Internet protocols and applications by reserving collections of geographically distributed hardware components and the network links connecting them, e.g., via Internet2 or NLR. A core concept for GENI and other cloud computing platforms is *resource leasing*.

Since GENI allocates resources from multiple providers, it uses one or more *Clearinghouses* to mediate the allocation. Providers delegate the right to allocate their resources to these Clearinghouses, which aggregate the resources and allocate them to researchers. As with Amazon's EC2 and EBS, GENI allocates virtualized hardware components to leverage statistical multiplexing and allow multiple researchers to use one hardware component simultaneously.

GENI consists of multiple consumers that acquire resources from one or more Clearinghouses that act as brokers for transactions between providers and consumers. The initial intent is for the GENI Project Office to operate a small number of Clearinghouses, but, in general, there may be multiple Clearinghouses operated by governments, companies, or university-led consortiums. While the initial prototype's scale does not warrant market-based allocation mechanisms, reaching GENI's goal for Internet-scale operation—allocating millions of components—motivates a market-oriented approach. Further, decentralizing resource allocation among multiple Clearinghouses gives GENI the flexibility to introduce market-oriented approaches incrementally in only a few Clearinghouses initially. We chose GENI as our motivation because its decentralized design is amenable to incrementally introducing market-oriented approaches and its structure is still open for debate. Further, we believe

GENI’s goal as a platform for experimental research should also include research on its own resource allocation mechanisms.

In a market-oriented GENI, consumers increase their utility by purchasing resources from Clearinghouses and satisfying their resource requirements. Clearinghouses allocate resources to maximize their profit—the difference between their revenue from consumers and their cost of providing resources. As motivation for a market-oriented approach, we also assume that the demand for GENI’s resources will exceed its supply, which has been the case for GENI’s primary predecessor Planet-Lab [19, 55]. In general, the market mechanism to determine the resource allocation could be either *centralized*, e.g., auction, or *distributed*, e.g., negotiation. Distributed approaches like negotiation, where allocations emerge as the result of a sequence of interactions between self-interested agents, are well-suited to GENI, since its scale and dynamics preclude a once-and-for-all global optimization of resource usage [1].

6.3 The Negotiation Model

6.3.1 The Resource Allocation Problem

We treat each consumer as a buyer and each provider as a seller, where \mathcal{B} denotes the set of buyers and \mathcal{S} denotes the set of sellers. Each buyer $\mathbf{b} \in \mathcal{B}$ has a high level task τ , such as an experiment. The task τ of buyer \mathbf{b} has the following attributes:

- A resource set $\mathcal{R}_{\mathbf{b}}$ and the quantity of units τ requires. For a resource $r \in \mathcal{R}_{\mathbf{b}}$, τ requires $q(\mathcal{R}_{\mathbf{b}}, r)$ units of resource r .
- Task generation time $tg(\mathbf{b})$ when the task is generated.
- Earliest start time $est(\mathbf{b})$ where task τ cannot start before time $est(\mathbf{b})$. Generally $est(\mathbf{b}) > tg(\mathbf{b})$ and \mathbf{b} can use the time between $est(\mathbf{b})$ and $tg(\mathbf{b})$ to acquire resources.

- The period $pd(\mathbf{b})$ of resource usage, such that \mathbf{b} must use resources $\mathcal{R}_{\mathbf{b}}$ for a period of length $pd(\mathbf{b})$.
- Deadline $dl(\mathbf{b})$ that indicates the latest start time of the task of a buyer \mathbf{b} . Since $dl(\mathbf{b}) \geq est(\mathbf{b})$. If $dl(\mathbf{b}) > est(\mathbf{b})$, the buyer has the flexibility to determine the start time of the experiment. Note that the task must finish before $dl(\mathbf{b}) + pd(\mathbf{b})$, and a rational buyer will not negotiate after $dl(\mathbf{b})$.
- Value $v_{\mathbf{b}}(t)$ represents the value \mathbf{b} attaches to task completion as a function of completion time t . Following [86], \mathbf{b} has its maximum value at time $est(\mathbf{b}) + pd(\mathbf{b})$ and its minimum value at time $dl(\mathbf{b}) + pd(\mathbf{b})$.

Each seller $\mathbf{s} \in \mathcal{S}$ has different types of resources $\mathcal{R}_{\mathbf{s}}$ in varying quantities, $q(\mathcal{R}_{\mathbf{s}}, r)$ units of resource $r \in \mathcal{R}_{\mathbf{s}}$, and suffers a cost $c_{\mathbf{s}}(r)$ for providing each unit of resource $r \in \mathcal{R}_{\mathbf{s}}$ for a unit time period. This model follows GENI in that sellers have different types of resources, although we simplify our problem by allowing only one “plan” for each task. While we specify only a single set of resources to satisfy each task, in general, multiple different types of resources may be able to satisfy a task. For example, a researcher may either plan an experiment with a small number of resources for a long duration, or a large number of resources and a short duration. In these cases, we can extend our formulation to include multiple plans.

We assume each buyer is able to discover the set of resources each seller provides. This assumption is reasonable since each seller is willing to let others to know its capability, and, from a single agent’s perspective, knowing other agents’ information may help it to develop appropriate strategies. For example, if a buyer knows that the resource competition is low, it may offer a lower price. We assume that 1) each buyer knows each seller’s expected cost $c_{\mathbf{b}}(r)$ of providing a resource r ; and 2) each agent has knowledge about the demand/supply ratio $\psi(r)$ of resource r over time. This assumption is not more restrictive than related work [11, 98]. Further, in dynamic

Table 6.1. Symbols used in this chapter

\mathcal{R}_b	the set of resources needed by buyer b
\mathcal{R}_s	the set of resources provided by seller s
$q(\mathcal{R}_b, r)$	the quantity of resource r needed by buyer b
$q(\mathcal{R}_s, r)$	the quantity of resource r provided by seller s
$tg(\mathbf{b})$	b 's task generation time
$est(\mathbf{b})$	earliest start time of b 's task
$pd(\mathbf{b})$	resource usage period of b 's task
$dl(\mathbf{b})$	deadline of b 's task
$v_b(t)$	b 's value of completing its task at time t
$c_s(r)$	s 's cost for providing one unit of r for a unit time period
$\psi(r)$	demand/supply ratio of resource r
$\mathcal{A}_b/\mathcal{A}_s$	b / s 's final agreements
$\mathcal{TA}_b/\mathcal{TA}_s$	b / s 's tentative agreement set
$\mathcal{A}_b^f/\mathcal{TA}_b^f$	the set of final/tentative full agreements
$R(\mathcal{A})$	the set of resources provided by the agreement set \mathcal{A}
\mathcal{RA}_s	s 's running agreements
\mathcal{KA}_s	s 's set of final agreements not to decommit
\mathcal{KTA}_s	s 's set of tentative agreements to not to cancel
\mathcal{AO}_s	s 's set of offers to accept

markets, a buyer can estimate a seller's cost and market competition by analyzing its negotiation history. We explore the sensitivity of this assumption in our experiments.

6.3.2 Negotiation Protocol

This work extends the alternating offers protocol [111], which has been widely used for bilateral bargaining. Before we formally define the protocol, we first define agents' possible actions:

- *offer*[o], where o is buyer **b**'s offer to a seller **s**. An offer o is of the form $\langle pr, pe, \mathcal{R}, est, pd, dl \rangle$ where pr is the offering price, pe is the decommitment penalty, \mathcal{R} specifies the set of required resources and their quantities, est is the

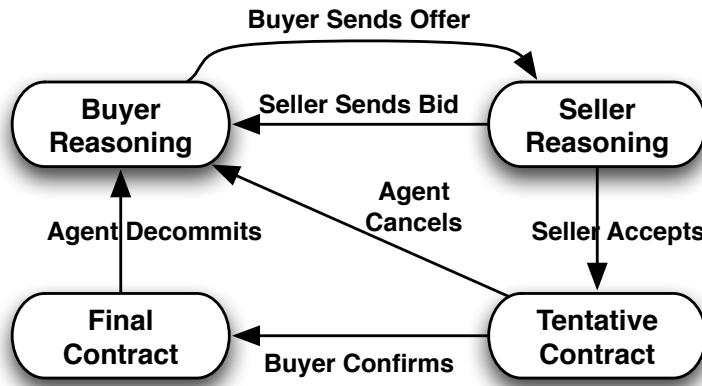


Figure 6.3. Finite state machine for the negotiation protocol

earliest start time of providing resources, pd is the duration, and dl is the latest time for providing resources. Note that when $dl > est$, the buyer provides a flexible schedule and the receiving seller is able to decide the exact start time of providing resources.

- *accept*[o]. When a seller s receives an offer o' , s is able to accept the offer resulting in the two agents reaching a *tentative* agreement. If $dl(o') > est(o')$, s must decide the exact start time of providing $\mathcal{R}(o')$ and $dl(o)$ should be equal to $est(o)$.
- *bid*[\mathcal{Q}]. When a seller s receives an offer o' and the offer is not acceptable, s can send quotes \mathcal{Q} for its available resources. Each quote $Q \in \mathcal{Q}$ describes the quantity of resource $r \in \mathcal{R}(o')$ and its asking price.
- *confirm*[o]. When a seller s accepts an offer o , two agents reach a *tentative* agreement and the buyer can *confirm* the tentative agreement. If b confirms the tentative agreement, then the agreement becomes a *final* agreement.

- *cancel*[*o*]. After two agents make a tentative agreement, any agent can *cancel* the agreement *without* paying a penalty. Then, negotiation between the two agents fails with no agreement.
- *decommit*[*o*]. After a final agreement is made, an agent has the opportunity to decommit from the agreement and the decommitting agent pays the penalty to the other party. Note that after time $est(o)$, no agent can decommit from the agreement. Furthermore, the decommitting agent pays the penalty after decommitment happens.

Figure 6.3 shows the finite state machine for the negotiation between **b** and **s**. The initial state is “*buyer reasoning*” in which **b** decides how to make the offer. After **b** sends an offer to **s**, the state is “*seller reasoning*” in which **s** is deciding whether to accept the offer or make a bid. If **s** accepts the offer, a tentative agreement is made. Otherwise, **s** sends a bid to **b** and then it is **b**’s turn to decide its offer. If **b** confirms a tentative agreement, the negotiation is in the “*final agreement*” state. If one agent cancels a tentative agreement or decommits from a final agreement, their negotiation fails and **b** can restart to make an offer.

An important feature of our negotiation model is that many buyer-seller pairs can negotiate simultaneously. In addition to the decommitting action, we also introduce another pair of actions “*confirm*” and “*cancel*”. With the two actions, if a seller accepts an offer, a buyer can still have the chance to “*decommit*” from the agreement without paying a penalty. Assume that **b** only needs one resource. In absence of the action *cancel*, if **b** makes offers to multiple sellers that all accept, **b** must buy multiple items or decommit from agreements by paying penalties. Accordingly, **b** may only propose to one seller. In presence of actions *cancel* and *confirm*, **b** can choose only one contract while negotiating with multiple sellers simultaneously.

Next we formalize the notion of utility. The utility of buyer **b** depends on its task completion time and its payment, including 1) its payment for getting resources, and

2) penalties it pays to other agents and receives from other agents. \mathbf{b} 's utility at time t is

$$u_{\mathbf{b}}(t) = \begin{cases} v_{\mathbf{b}}(t) + \rho_{\mathbf{b}} & \text{if } \mathbf{b}'\text{s task is finished} \\ \rho_{\mathbf{b}} & \text{otherwise} \end{cases}$$

where $\rho_{\mathbf{b}}$ is the balance of the buyer \mathbf{b} —the difference between the payment received and the payment paid to other agents.

The total utility of each seller $\mathbf{s} \in \mathcal{S}$ from time 0 to time t is $u_{\mathbf{s}}(t) = \rho_{\mathbf{s}} - c_{\mathbf{s}}$ where $\rho_{\mathbf{s}}$ is the balance of the seller \mathbf{s} at time t and $c_{\mathbf{s}}$ is seller \mathbf{s} 's cost for providing resources from the beginning to time t .

6.4 Buyers' Negotiation Strategy

Before formally defining a buyer \mathbf{b} 's negotiation strategy, we first discuss other important factors we consider:

- **Deadline Pressure.** \mathbf{b} must satisfy its resource requirements by the deadline $dl(\mathbf{b})$, which is a hard constraint.
- **Sellers' Cost.** A rational seller will not accept a price lower than its cost. A buyer needs to offer different prices for different resources which have different costs.
- **Single Provider.** If $\mathcal{R}_{\mathbf{b}} \subseteq \mathcal{R}_{\mathbf{s}}$, \mathbf{b} can make a *full* agreement with a single seller \mathbf{s} which can satisfy \mathbf{b} 's resource requirements. Otherwise, it must request resources from different sellers and make a set of *partial* agreements, each of which can only satisfy part of \mathbf{b} 's resource requirements. The negotiation for the latter case is more complex since \mathbf{b} must have contracts with multiple sellers, and making no agreement may be better than making agreements which cannot satisfy \mathbf{b} 's requirements. Furthermore, if \mathbf{b} 's resource requirements are satisfied through a set of contracts, the set of contracts should be *compatible* in that all contracts should provide resources during the same time frame.

In summary, a buyer agent's optimal action at each time point is affected by many factors and it is impossible to construct an integrated framework in which all these factors are optimized concurrently. Instead, this work connects those inter-dependent factors indirectly and develops a set of heuristics to approximate agents' decision making. In what follows we first introduce buyer \mathbf{b} 's strategy (Algorithm 6) informally and then present it formally.

Algorithm 6: Negotiation strategy of buyer \mathbf{b}

```

Set  $\mathcal{A}_{\mathbf{b}} = \emptyset$ ,  $\mathcal{TA}_{\mathbf{b}} = \emptyset$ ,  $t$  is the real time (initially,  $t = tg(\mathbf{b})$ ).
Let  $estp = est(\mathbf{b})$ ,  $dlp = dl(\mathbf{b})$  be the earliest start time and deadline for negotiating partial contracts.
Let  $\kappa$  (e.g., 4) be the total number of times to try different execution schedules of partial agreements. Let
 $T_{bk} = (dl(\mathbf{b}) - tg(\mathbf{b}))/\kappa + tg(\mathbf{b})$ .
while  $t < dl(\mathbf{b})$  and the task has not started to run do                                     /* main loop */
|   foreach  $s \in \mathcal{S}$  such that  $\mathcal{R}_s \cap \mathcal{R}_{\mathbf{b}} \neq \emptyset$  and  $\mathbf{b}$  is not negotiating with  $s$  do
|   |   send offer GENERATE_OFFER( $\mathcal{A}_{\mathbf{b}}, \mathcal{TA}_{\mathbf{b}}, s$ ) to seller  $s$ ;
|   end
|   if seller  $s$  sends a bid  $\mathcal{Q}$  then
|   |   update the bid set;
|   end
|   if seller  $s$  accepts offer  $o$  then
|   |   EVALUATE_ACCEPT( $\mathcal{A}_{\mathbf{b}}, \mathcal{TA}_{\mathbf{b}}, o, s$ );
|   end
|   if  $t > estp$  then
|   |   decommit (cancel) agreements  $\mathcal{A}_{\mathbf{b}} - \mathcal{A}_{\mathbf{b}}^f$  ( $\mathcal{TA}_{\mathbf{b}} - \mathcal{TA}_{\mathbf{b}}^f$ );
|   |   set  $estp = \max\{t, est(\mathbf{b})\}$ ,  $dlp = dl(\mathbf{b})$ ;
|   end
|   if  $t \geq T_{bk}$  then
|   |    $\kappa - \kappa$ ,  $T_{bk} = (dl(\mathbf{b}) - t)/\kappa + t$ ;
|   |   if  $R(\mathcal{A}_{\mathbf{b}} + \mathcal{TA}_{\mathbf{b}} - \mathcal{A}_{\mathbf{b}}^f - \mathcal{TA}_{\mathbf{b}}^f) \subset \mathcal{R}_{\mathbf{b}}$  then
|   |   |   decommit (cancel) agreements  $\mathcal{A}_{\mathbf{b}} - \mathcal{A}_{\mathbf{b}}^f$  ( $\mathcal{TA}_{\mathbf{b}} - \mathcal{TA}_{\mathbf{b}}^f$ );
|   |   |   set  $estp = \max\{t, est(\mathbf{b})\}$ ,  $dlp = dl(\mathbf{b})$ ;
|   |   end
|   end
|   if seller  $s$  decommits from agreement  $o$  then
|   |   remove  $o$  from  $\mathcal{A}_{\mathbf{b}}$ ;
|   end
|   if seller  $s$  cancel tentative agreement  $o$  then
|   |   remove  $o$  from  $\mathcal{TA}_{\mathbf{b}}$ ;
|   end
|   if seller  $s$  has not responded to  $\mathbf{b}$ 's proposing  $o$  for a period  $\epsilon$  then
|   |   send offer GENERATE_OFFER( $\mathcal{A}_{\mathbf{b}}, \mathcal{TA}_{\mathbf{b}}, s$ ) to seller  $s$ ;
|   end
|   if seller  $s$  has not responded to  $\mathbf{b}$ 's accepting offer  $o$  for a period  $\epsilon$  then
|   |   cancel the tentative agreement  $o$  and remove  $o$  from  $\mathcal{TA}_{\mathbf{b}}$ ;
|   end
end
cancel from all tentative agreements  $\mathcal{TA}_{\mathbf{b}}$ ;
if the task has started to run then
|   decommit from each agreement  $o \in \mathcal{A}_{\mathbf{b}}$  if  $o$  is useless and  $pr(o) > pe(o)$ ;
else
|   decommit from each agreement  $o \in \mathcal{A}_{\mathbf{b}}$  if  $pr(o) > pe(o)$ ;
end

```

One distinguishing feature of \mathbf{b} 's negotiation strategy is that it always tries to make two sets of agreements both of which can satisfy its resource requirements.

Therefore, if a set of agreements is decommitted, \mathbf{b} can use the other agreement set to satisfy its resource requirements. If both set of agreements are not decommitted, when one set of agreements starts execution, \mathbf{b} can decommit from the other set of agreements. If the start time of two sets of agreements are the same, \mathbf{b} will choose one set of agreements to decommit before the execution starts. Specifically, \mathbf{b} is always trying to make a final full agreement and a set of partial final agreements both of which can satisfy its resource requirements. In case no single seller can satisfy \mathbf{b} 's resource requirements, \mathbf{b} makes two sets of partial final agreements. In addition, \mathbf{b} sets a small penalty for each partial agreement and thus it only needs to pay a small penalty for decommitting from any partial agreement. While a buyer can make more agreements to increase the probability that its task can be finished, it has to pay more for those agreements since for each unnecessary agreement, it has to pay either the penalty or the agreement price. Alternatively, if a buyer only makes one set of agreements, it may be difficult to find another set of agreements to satisfy the buyer's resource requirements when some agreements are decommitted. Experimental results show that making two sets of agreements is better than making only one set of agreements and making more than two sets of agreements. While the cloud resource allocation problem is different from the problem in the previous chapter, our results correspond to our findings in the previous chapter that a buyer making two or three tentative agreements always gained the highest utility.

Another distinguishing feature of \mathbf{b} 's negotiation strategy is that while deciding the offering price $pr(\mathcal{R}, est, dl, t)$ of requesting resources \mathcal{R} at time t with earliest start time est and latest start time dl , the following factors are considered. First, the pressure of deadline. The buyer makes more concessions when the deadline approaches. Such time-dependent concession strategies have been widely used in the literature [11, 48]. Second, the cost $c_b(r)$ of resource r . Intuitively, a buyer needs to pay more for a resource with a higher cost. Third, the demand/supply ratio $\psi(r)$

Algorithm 7: GENERATE_OFFER($\mathcal{A}_b, \mathcal{TA}_b, s$)

```

if  $|\mathcal{A}_b^f \cup \mathcal{TA}_b^f| = 0$  and  $\mathcal{R}_b \subseteq \mathcal{R}_s$  then
  Let price be  $pr = pr(\mathcal{R}_b, \max\{t, est(\mathbf{b})\}, dl(\mathbf{b}), t)$  using Eq. (6.1);
  Let the penalty be  $pe = v_b(\max\{t, est(\mathbf{b})\} + pd(\mathbf{b})) - pr$ ;
  return offer  $o = \langle pr, pe, \mathcal{R}_b, \max\{t, est(\mathbf{b})\}, pd(\mathbf{b}), dl(\mathbf{b}) \rangle$ ;
else
  Let  $\mathcal{R} = \mathcal{R}_s \cap (\mathcal{R}_b - R(\mathcal{A}_b + \mathcal{TA}_b - \mathcal{A}_b^f - \mathcal{TA}_b^f))$ ;
  if  $\mathcal{R} = \emptyset$  then
    return offer  $o = null$ ;
  end
  if  $estp \neq dlp$  then
    If possible, set the value of  $estp = dlp > t + \sigma$  based on bids from sellers such that the available
    resource from  $estp$  to  $estp + pd(\mathbf{b})$  can satisfy  $\mathbf{b}$ 's resource requirements;
  end
  Let price be  $pr = pr(\mathcal{R}, estp, dlp, t)$  using Eq. (6.2);
  Let the penalty be  $pe = \alpha \cdot pr$  (e.g.,  $\alpha = 0.05$ );
  return offer  $o = \langle pr, pe, \mathcal{R}, estp, pd(\mathbf{b}), dlp \rangle$ ;
end

```

Algorithm 8: EVALUATE_ACCEPT($\mathcal{A}_b, \mathcal{TA}_b, o, s$)

```

Let new offer  $o' = \text{GENERATE\_OFFER}(\mathcal{A}_b, \mathcal{TA}_b, s)$ ;
if 1)  $\mathcal{R}(o) = \mathcal{R}(o') = \mathcal{R}_b$ ,  $|\mathcal{A}_b^f \cup \mathcal{TA}_b^f| = 0$ , and  $pr(o) \leq pr(o')$  or 2)  $\mathcal{R}(o) = \mathcal{R}(o')$ ,
 $est(o) = est(o') = estp = dlp$ , and  $pr(o) \leq pr(o')$  then
  confirm agreement  $o$  and add it to  $\mathcal{A}_b$ ;
else
  cancel agreement  $o$ ;
end

```

of a resource r . The higher the ratio, the higher the price for the resource. Market (resource) competition has the largest effect on the equilibrium price [7]. Formally, the offering price $pr(\mathcal{R}_b, est, dl, t)$ for all resources \mathcal{R}_b is defined as

$$c(\mathcal{R}_b) + (RP(est, dl) - c(\mathcal{R}_b)) \left(\frac{t - tg(\mathbf{b})}{dl(\mathbf{b}) - tg(\mathbf{b})} \right)^\varepsilon \quad (6.1)$$

where $c(\mathcal{R}_b) = \sum_{r \in \mathcal{R}_b} c_b(r) q(\mathcal{R}_b, r) pd(\mathbf{b})$ is the expected cost of resources \mathcal{R}_b and $RP(est, dl)$ is the expected value of finishing the task with earliest start time est and latest start time dl . Formally,

$$RP(est, dl) = \begin{cases} \frac{\int_{est}^{dl} v_b(pd(\mathbf{b}) + t') dt'}{dl - est} & \text{if } dl \neq est \\ v_b(pd(\mathbf{b}) + est) & \text{otherwise} \end{cases}$$

When $t = tg(\mathbf{b})$, $pr(\mathcal{R}_b, est, dl, t) = c(\mathcal{R}_b)$, which is the lowest offer acceptable to sellers. When $t = dl(\mathbf{b})$, $pr(\mathcal{R}_b, est, dl, t) = RP(est, dl)$, which is the highest offering price of \mathbf{b} since the buyer will get negative utility if it pays more than its

value of finishing the task. Parameter $\varepsilon > 0$ is used to model how the buyer \mathbf{b} increases its offering price with the increase of time t . With infinitely many values of ε , there are infinitely many possible strategies in making concessions with respect to the remaining time. However, they can be classified into: 1) *Linear*: $\varepsilon = 1$, 2) *Conciliatory*: $0 < \varepsilon < 1$, and 3) *Conservative*: $\varepsilon > 1$ [48]. We adopt the linear strategy for \mathbf{b} .

By considering both resources' costs and market competition, the buyer's offering price for $\mathcal{R} \subset \mathcal{R}_{\mathbf{b}}$ is calculated in the following way:

$$pr(\mathcal{R}, est, dl, t) = \sum_{r \in \mathcal{R}} q(\mathcal{R}, r) pd(\mathbf{b}) pr(r, est, dl, t) \quad (6.2)$$

$$pr(r, est, dl, t) = c_{\mathbf{b}}(r) + \frac{(pr(\mathcal{R}_{\mathbf{b}}, est, dl, t) - c(\mathcal{R}_{\mathbf{b}})) \psi(r) c_{\mathbf{b}}(r)}{pd(\mathbf{b}) \sum_{r \in \mathcal{R}_{\mathbf{b}}} \psi(r) c_{\mathbf{b}}(r) q(\mathcal{R}_{\mathbf{b}}, r)}$$

$pr(r, est, dl, t)$ is the price for one unit of resource r and it increases with its cost $c_{\mathbf{b}}(r)$ and the demand/supply ratio $\psi(r)$.

Let $\mathcal{A}_{\mathbf{b}}$ be \mathbf{b} 's final agreements and $\mathcal{TA}_{\mathbf{b}}$ be \mathbf{b} 's tentative agreement set. Let $\mathcal{A}_{\mathbf{b}}^f \subseteq \mathcal{A}_{\mathbf{b}}$ ($\mathcal{TA}_{\mathbf{b}}^f \subseteq \mathcal{TA}_{\mathbf{b}}$, respectively) be the set of final full (tentative, respectively) agreements. Let $R(\mathcal{A})$ be the set of resources provided by the agreement set \mathcal{A} . If \mathbf{b} has no full agreement, i.e., $|\mathcal{A}_{\mathbf{b}}^f \cup \mathcal{TA}_{\mathbf{b}}^f| = 0$, it will request for all resources $\mathcal{R}_{\mathbf{b}}$ from sellers which can satisfy its full resource requirements and request for resources $\mathcal{R}_{\mathbf{s}} \cap (\mathcal{R}_{\mathbf{b}} - R(\mathcal{A}_{\mathbf{b}} + \mathcal{TA}_{\mathbf{b}} - \mathcal{A}_{\mathbf{b}}^f - \mathcal{TA}_{\mathbf{b}}^f))$ from each seller \mathbf{s} which can only satisfy part of its resource requirements. If \mathbf{b} has a full agreement, it will request for resources $\mathcal{R}_{\mathbf{s}} \cap (\mathcal{R}_{\mathbf{b}} - R(\mathcal{A}_{\mathbf{b}} + \mathcal{TA}_{\mathbf{b}} - \mathcal{A}_{\mathbf{b}}^f - \mathcal{TA}_{\mathbf{b}}^f))$ from each seller $\mathbf{s} \in \mathcal{S}$.

When buyer \mathbf{b} wants to acquire resources \mathcal{R} from seller \mathbf{s} at time t , in addition to specifying the offering price, it also decides the decommitment penalty pe , and the task execution period. First consider the case in which $|\mathcal{A}_{\mathbf{b}}^f \cup \mathcal{TA}_{\mathbf{b}}^f| = 0$ and $\mathcal{R} = \mathcal{R}_{\mathbf{b}}$. In this case, \mathbf{b} simply requests its earliest execution start time $est(\mathbf{b})$,

deadline $dl(\mathbf{b})$, and the execution period $dl(\mathbf{b})$. The seller will decide the exact start time. We use a simple rule to decide the decommitment penalty: the lower the price $pr(\mathcal{R}_{\mathbf{b}}, est, dl, t)$, the higher the penalty. In other words, \mathbf{b} does not want a cheap full agreement to be decommitted. One example rule to set the penalty is $pe = v_{\mathbf{b}}(\max\{t, est(\mathbf{b})\} + pd(\mathbf{b})) - pr$.

We also consider the case $|\mathcal{A}_{\mathbf{b}}^f \cup \mathcal{T}\mathcal{A}_{\mathbf{b}}^f| > 0$ or $\mathcal{R} \neq \mathcal{R}_{\mathbf{b}}$. In this case, \mathbf{b} must decide what time period to request resources since different sellers need to provide resources in the same time period and this decision making is difficult due to uncertainty and agents' selfishness. In this work, \mathbf{b} decides the task execution schedule for partial agreements based on its information about sellers' available resources, which can be obtained from the bid messages and acceptance messages from sellers. Note that there is no guarantee that \mathbf{b} can get part or all of \mathbf{s} 's available resources due to the market dynamics. Specifically, \mathbf{b} searches from time $\max\{t + \sigma, est(\mathbf{b})\}$ until $dl(\mathbf{b})$ and sets the task start time est as the earliest time point from which sellers' available resources from time est to $est + pd(\mathbf{b})$ can satisfy the buyer's resource requirements. We use the parameter $\sigma > 0$ to allow the buyer the flexibility to negotiate for resources. We choose this simple rule for two reasons. First, since a buyer's value of finishing a task generally decreases with the task start time, the buyer can potentially achieve a higher utility if it negotiates for a set of agreements with an early task start time. Second, due to market dynamics and agents' strategic interaction, it is impossible to determine the best start time. If there is no start time for which the buyer's resource requirements can be satisfied, the buyer simply sets $est = est(\mathbf{b})$ and $dl = dl(\mathbf{b})$ and it will not confirm any partial agreement. Using our simple rule, we set the decommitment penalty in this case to $pe = \alpha \cdot pr$, where $0 < \alpha < 0.2$.

Once the task execution schedule of partial agreements is determined, buyer \mathbf{b} will request resources from sellers according to the task execution schedule. However, the selected task execution schedule may cause buyer \mathbf{b} to fail to find agreements

to satisfy its resource requirements. Therefore, it is important for buyer **b** to try other task execution schedules if it fails to get agreements with the current schedule. In other words, buyer **b** should have the “backtracking” ability of changing its task execution schedule. In this work, a buyer agent will change its task execution schedule if it fails to satisfy its resource requirements for a given time. When a buyer **b** changes its task execution schedule, it will first decommit from its other partial agreements.

6.5 Sellers’ Negotiation Strategy

Algorithm 9: Negotiation strategy of seller **s**

```

Set  $\mathcal{A}_s = \emptyset, \mathcal{TA}_s = \emptyset;$ 
if buyer b decommits from agreement  $o$  then
  | remove  $o$  from  $\mathcal{A}_b;$ 
end
if buyer b cancels tentative agreement  $o$  then
  | remove  $o$  from  $\mathcal{TA}_b;$ 
end
if buyer b confirms tentative agreement  $o$  then
  | remove  $o$  from  $\mathcal{TA}_b$  and add  $o$  to  $\mathcal{A}_b;$ 
end
if buyer b sends an offer  $o$  then
  | run the greedy algorithm for  $\text{OPT}_s(\mathcal{RA}_s, \mathcal{A}_s, \mathcal{TA}_s, \mathcal{O}_s);$ 
end
if buyer b has not responded to s's proposing  $o$  for a period  $\epsilon$  then
  | send offer to buyer s with a price  $c_s(r)\varphi(r, t)$  for each resource  $r;$ 
end
if buyer b has not responded to s's accepting offer  $o$  for a period  $\epsilon$  then
  | cancel the tentative agreement  $o$  and remove  $o$  from  $\mathcal{TA}_s;$ 
end

```

Our negotiation strategy for the seller (Algorithm 9) has two features. First, the seller adopts a “myopic” negotiation strategy in the sense that it accepts an offer if and only if it can gain some *immediate* payoff by accepting the offer, and will not consider the effect of its current action on the future utilities. Part of the reason is that the seller has limited information about other agents in the market and it is impractical to make assumptions about behavior of other selfish agents in the market. In addition, when a seller receives an offer, it will first make acceptance and decommitment decisions, and then generate bids to the buyer if the offer is not acceptable. Second, the seller decides the acceptable price for a set of resources based on resource competition and cost of resources.

If the competition of a resource is high, a seller has an expectation that it will receive a high price for the resource. When a seller receives an offer o , it first generates a threshold price $\phi(o)$. If $pr(o) < \phi(o)$, it will not accept the offer. The threshold price $\phi(o)$ is defined as

$$\phi(o) = \sum_{r \in \mathcal{R}(o)} q(\mathcal{R}(o), r) c_s(r) (1 + \psi(r)) pd(o)$$

in which $c_s(r)(1 + \psi(r))pd(o)$ is seller s 's "asking" price for one unit of resource r . Obviously, $\phi(o) > cost(o) = \sum_{r \in \mathcal{R}(o)} q(\mathcal{R}(o), r) c_s(r) pd(o)$. If an offer is not acceptable, the seller will simply report its available resources in the buyer's request, as well as the unit price of each resource r as $c_s(r)(1 + \psi(r))$.

Since an agent can decommit from a final agreement, a seller can make more agreements than its capacity. In this case, a seller can decommit from an unsatisfiable agreement before the resource providing time. However, since an agent does not know whether the other agent will decommit from an agreement and the seller may pay a high penalty, we have chosen a seller's strategy where the seller may not make agreements beyond its capability. That is, without decommitting from any final agreement or canceling any tentative agreement, the seller must be able to fulfill its current running agreements \mathcal{RA}_s , final agreements \mathcal{A}_s , and tentative agreements \mathcal{TA}_s [115]. This strategy also implies that when a seller receives a message indicating confirmation of an agreement, it can fulfill the agreement without decommitting from any final agreement or cancel any tentative agreement.

The most difficult decision problem for the seller is how to handle a set of acceptable offers, which can be formulated as an optimization problem $\mathbf{OPT}_s(\mathcal{RA}_s, \mathcal{A}_s, \mathcal{TA}_s, \mathcal{O}_s)$: Given the running agreements \mathcal{RA}_s , final agreements \mathcal{A}_s , tentative agreements \mathcal{TA}_s , and offers \mathcal{O}_s , compute the set \mathcal{KA}_s of final agreements not to decommit, the set \mathcal{KTA}_s of tentative agreements to not to cancel, and the set \mathcal{AO}_s of offers to accept to maximize the following objective function

$$\sum_{o \in \mathcal{KA}_s \cup \mathcal{KA}_s \cup \mathcal{AO}_s} pr(o) - cost(o) - \sum_{o \in \mathcal{A}_s - \mathcal{KA}_s} pe(o)$$

with the constraint that \mathbf{s} can fulfill final agreements \mathcal{KA}_s and tentative agreements $\mathcal{TA}_s \cup \mathcal{AO}_s$.

Theorem 27. *The optimization problem $\text{OPT}_s(\mathcal{RA}_s, \mathcal{A}_s, \mathcal{TA}_s, \mathcal{O}_s)$ is \mathcal{NP} -complete.*

The theorem's proof is a straightforward reduction from the 0-1 Knapsack problem. Thus, we propose a greedy algorithm to handle this computationally costly optimization problem. First, an agreement is treated as an offer and let $\Omega = \mathcal{A}_s \cup \mathcal{TA}_s \cup \mathcal{O}_s$ be the set of offers that must be considered. \mathbf{s} 's revenue of accepting offer $o \in \Omega$ is

$$rv(o) = \begin{cases} pr(o) - cost(o) + pe(o) & \text{if } o \in \mathcal{A}_s \\ pr(o) - cost(o) & \text{otherwise} \end{cases}$$

Next, all the offers Ω are sorted by decreasing revenue and offers are greedily picked in this order, starting with the first offer, and until no offers remain. Let $\Omega' = \emptyset$ be the set of accepted offers. When an offer o is picked, add o to Ω' and check whether the seller is able to fulfill all agreements Ω' . Note that if $o \in \mathcal{O}_s$ and $dl(o) > est(o)$, the seller must decide the schedule (the start time of providing resources and end time of providing resources) for providing resources specified in the offer. If the seller can fulfill all the agreements in Ω' and o is an offer from a buyer, the seller will send an acceptance message to the buyer. If the seller cannot fulfill all the agreements in Ω' , remove o from Ω' . If $o \in \mathcal{A}_s$, then send a decommitment message to the buyer involved in the agreement and pay the penalty. If $o \in \mathcal{TA}_s$, send a cancel message to the buyer involved in the agreement.

6.6 Empirical evaluation

To evaluate the performance of our mechanism, we implement a simulation testbed consisting of a discrete time virtual marketplace and a set of trading agents. We

Table 6.2. Variables

<i>Variables</i>	Values
<i>Number of sellers</i>	[5, 20]
<i>No. of resource types per seller</i>	[2, 8]
<i>Quantity of a resource per seller</i>	[2, 20]
<i>Unit cost of a resource</i>	[10, 100]
<i>No. of resource types per buyer</i>	[2, 6]
<i>Quantity of a resource per buyer</i>	[2, 8]
<i>Value/cost ratio</i>	[1.2, 5]
$pd(\mathbf{b})$	[10, 50]
$\frac{dl(\mathbf{b})-pd(\mathbf{b})-est(\mathbf{b})+1}{pd(\mathbf{b})}$ (<i>task execution flexibility</i>)	[0, 7]
$\frac{est(\mathbf{b})-tg(\mathbf{b})+1}{pd(\mathbf{b})}$ (<i>negotiation time ratio</i>)	[1, 8]
<i>resource demand/supply ratio $\psi(r)$</i>	[0.2, 10]

generate all seller agents before the market opens and buyers dynamically enter the market, which matches real-world environments with a fixed number of well-known sellers.

6.6.1 Different Mechanisms

Negotiation mechanism (NG): When a buyer enters the market, it negotiates with sellers following the protocol described in Section 6.3.2. At each time point, first all buyers are triggered and then all sellers are triggered. All agents employ the negotiation strategies described in Sections 6.4 and 6.5. A buyer quits the market when its task is finished or it fails to satisfy its resource requirements by its deadline. For comparison, we also implemented two other widely used mechanisms:

- *Combinatorial reverse auction (CRA)*: In combinatorial auctions [42], a large number of items are auctioned concurrently. In combinatorial reverse auctions, a buyer buys goods from many competing sellers. When a buyer enters the market, it announces its resource requirements, and sellers submit bids indicating the set of resources and their prices. Finally, the buyer determines the set of contracts. The buyer uses the well-known strategy-proof Vickrey auction

mechanism. We assume that each seller has no knowledge of other buyers and, thus, each seller truthfully reports its available resources and their costs.

- *Fixed price scheme (Amazon) [3]*: Amazon EC2 is a web service that provides resizable compute capacity. The primary pricing mechanism for Amazon is a fixed price scheme with hourly charges per virtual machine. While using the Amazon scheme, a seller sets its price for each resource in advance and sellers constantly update their available resources. When a buyer enters the market, it decides the set of resources to buy. In our experiments, we tried different methods for setting price of each resource, where the price/cost ratio is 1, 2, 3, or 5.

The Amazon scheme is similar to *CRA*, except that in the Amazon scheme, a seller's payment from buyers is decided by the seller. In contrast, a seller's payment in *CRA* is the opportunity cost that its presence introduces to all the other agents. Note that when the price/cost in the Amazon scheme is 1, the Amazon scheme is equivalent to *CRA* in terms of the allocation since each seller will only charge its cost. Our negotiation model is also similar to *CRA* since sellers' accepting offers in the negotiation model are equivalent to submitting bids in the auction model. There are two main differences between our negotiation model and the other two models. First, in the negotiation model, agents are allowed to decommit from agreements. Second, there is a dynamic bargaining process in the negotiation model.

6.6.2 Experimental Settings and Measures

We performed a series of experiments in a variety of test environments using the parameters from Table 6.2. The parameters are inspired by the current design of the GENI infrastructure [1]. In the experiments, the number of sellers are in the range of [5, 20], where each seller can provide 4 to 8 different types of resources. The quantity of a resource a seller can provide is in the range of [2, 20]. The cost of a resource per unit




Figure 6.4. Social welfare and resource competition

time is in the range of $[10, 100]$. Each buyer needs 2 to 6 different types of resources, and for each type of resource, a buyer needs 2 to 6 units. The length of resource usage is in the range of $[10, 50]$. The ratio $\frac{dl(\mathbf{b})-pd(\mathbf{b})-est(\mathbf{b})+1}{pd(\mathbf{b})} \in [0, 7]$ describes a buyer's flexibility of deciding when to start its task. Similarly, ratio $\frac{est(\mathbf{b})-tg(\mathbf{b})+1}{pd(\mathbf{b})} \in [1, 8]$ represents a buyer's time to negotiate for resources. We assume that each buyer has a linear value function in which the buyer gets the highest value when the task starts from $est(\mathbf{b})$ and the buyer gets the lowest value when the task starts at $dl(\mathbf{b})$. Value/cost ratio is used to generate a buyer's maximum value and minimum value based on sellers' cost of providing resources. $\psi(r) \in [0.2, 10]$ is the ratio of total resource requirements to total resource supply through the whole experiment horizon.

The main performance measure is the social welfare—the sum of all agents' utilities. Since the social welfare of a mechanism in different settings could be significantly different, we report the ratio of the social welfare of *CRA* and the Amazon mechanism to the social welfare of *NG*. We also report the success rate of different mechanisms—




Figure 6.5. Success rate and resource competition

the percentage of buyers which successfully complete their tasks. Note that a high success rate does not imply a high social welfare.

6.6.3 Results

Extensive stochastic simulations were carried out for all the combinations of variables in Table 6.2. For each combination, we randomly generated over 5000 experiments and for each experiment, and tried all the three mechanisms and generated average performance measures. Even though extensive stochastic simulations were carried out for all the situations, due to space limitations, we only present the representative results. The length of each experiment is 1000 time units. We found that the confidence interval for each average value is very tight around the value, so the confidence intervals are not reported.

6.6.3.1 Performance of the negotiation mechanism

Observation 1: NG achieved about 13% higher social welfare than any other evaluated mechanism. Figure 6.4 shows how the social welfare of different mechanisms

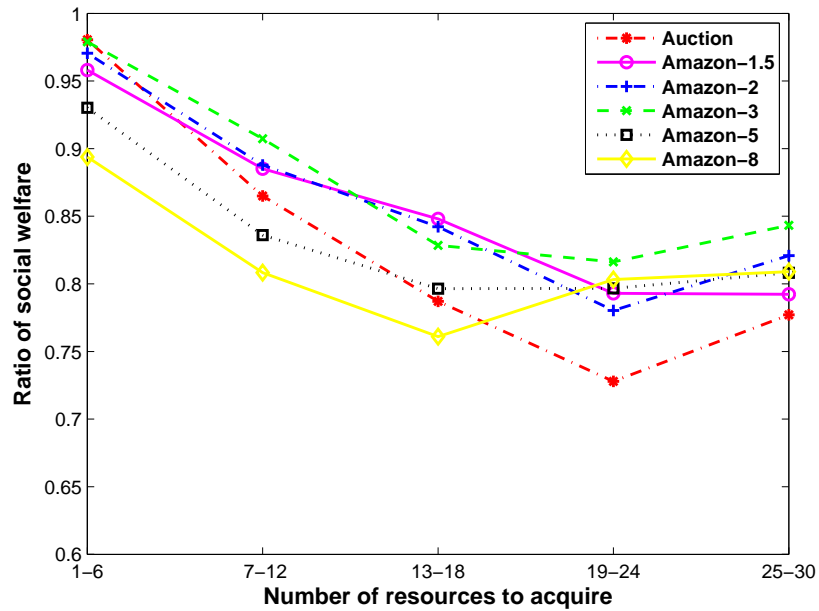


Figure 6.6. Social welfare and number of resource to acquire

changes with resource demand/supply ratio $\psi(r)$. We can observe that in all situations, NG 's social welfare is always higher than any other mechanism. Furthermore, when $\psi(r)$ is small (e.g., 0.2), CRA or the Amazon scheme with lower prices (e.g., Amazon-1.5) achieved higher social welfare than with higher prices (e.g., Amazon-8). In contrast, when $\psi(r)$ is large (e.g., 6), the Amazon scheme with higher prices (e.g., Amazon-8) achieved higher social welfare than CRA or Amazon scheme with lower prices. This observation is intuitive: When the resource competition is low, there are plenty of resources and each buyer can find them. However, when the resource competition is high, a mechanism can achieve a high social welfare if tasks with high revenues can be completed. If the price of each resource is low, a task with low revenue may get resources and a task with high revenue may fail to get resources since the resource were prematurely committed to the low revenue buyer and there was no way to decommit from the decision. In contrast, if a high price is set for each resource, only tasks with high revenues can get resources.




Figure 6.7. Success rate and number of resource to acquire

Figure 6.5 shows how the success rates of different mechanisms change with resource demand/supply ratio. First, a mechanism with a higher price has a lower success rate than that of a mechanism with a lower price. *NG*'s success rate is lower than some mechanisms with lower prices due to fact that in negotiation, each agent will not accept or offer any offer worse than its expectation. However, *NG*'s success rate is lower than that of any other mechanisms by no more than 10% when the resource demand/supply ratio is low and is almost the highest when the resource demand/supply ratio is higher than 1. Second, with the increase of resource competition, the success rate of each mechanism decreases, which corresponds to the intuition that with higher resource competition, it is more difficult to acquire resources.

Observation 2: Figure 6.6 shows how the social welfare of different mechanisms changes with the average number of resources acquired by buyers, which is $\sum_{r \in \mathcal{R}_b} q(\mathcal{R}_b, r)$. We can observe that the advantage of *NG* over other mechanisms increases with the number of resources to acquire. Figure 6.7 shows that the success




Figure 6.8. Social welfare and the flexibility of starting a task

rate decreases with the number of resources to acquire, which is intuitive since it is difficult to acquire more resources which have to be provided during the same period.

Observation 3: In some cases, the difference between a deadline and the earliest start time is large and each buyer has more flexibility of deciding when to start its task. A buyer \mathbf{b} can use the time between $est(\mathbf{b})$ and $dl(\mathbf{b})$ to negotiate for resources. As shown in Figure 6.9, the success rate of NG increases when buyers have more flexibility to decide when to start task execution. However, an agreement’s probability of being decommitted increases with more flexibility. Accordingly, a buyer may fail to get resources due to the decommitment. Figure 6.8 shows that, with the increase of the flexibility, the advantage of NG over the other mechanisms increases at the beginning and slightly decreases when buyers have a lot of flexibility to decide when to start task execution, which is mainly due to sellers’ decommitment.

Observation 4: A buyer \mathbf{b} can start negotiation at time $tg(\mathbf{b})$ and its task cannot start before $est(\mathbf{b})$. Figure 6.11 shows that NG ’s success rate increases with $(est(\mathbf{b}) - tg(\mathbf{b}))/pd(\mathbf{b})$ since a buyer has more time to negotiate for resources. However, as

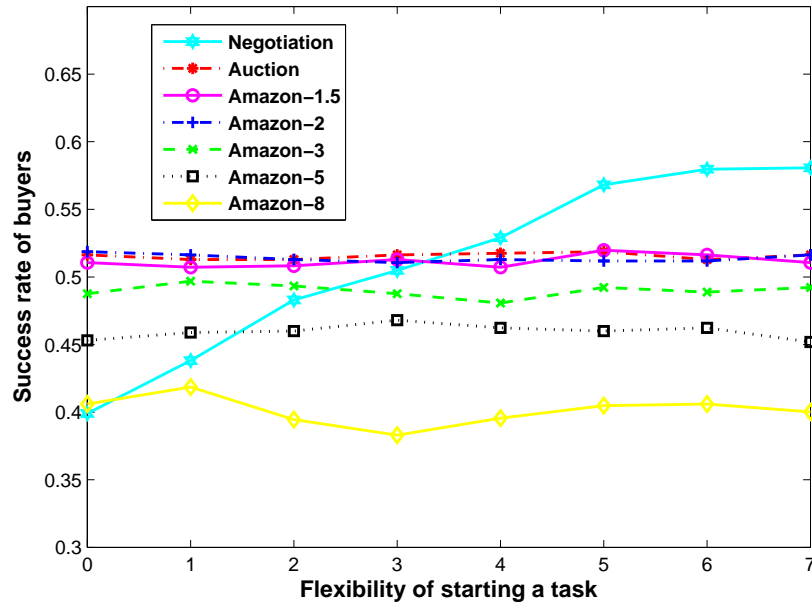


Figure 6.9. Success rate and the flexibility of starting a task

shown in Figure 6.10, the advantage of *NG* does not strictly increase with negotiation time: its advantage decreases when buyers have a long negotiation time. The reason is that a buyer’s agreements made at an early stage may be decommitted by sellers when there is a long negotiation deadline.

Observation 5: In addition to a fully distributed auction (*CRA*), we also designed a super buyer which receives requests from buyers and buys resources for buyers. The super buyer runs the auction when it has received a certain number of requests or one requesting buyer’s deadline is approaching, whichever occurs first. Experimental results show that *NG* still beat the centralized *CRA* by 11%. The centralized *CRA* beat the distributed *CRA* by no more than 2%.

One major difference between the distributed auction model and our negotiation mechanism is that in the negotiation mechanism agents are allowed to decommit from existing contracts at the cost of paying penalties. For comparison reason, we also allow agents to decommit in the auction model where decommitment penalties are set exogenously [11, 98], e.g., fixed penalties (e.g., {0, 10, 20, 40}) or penalty as a

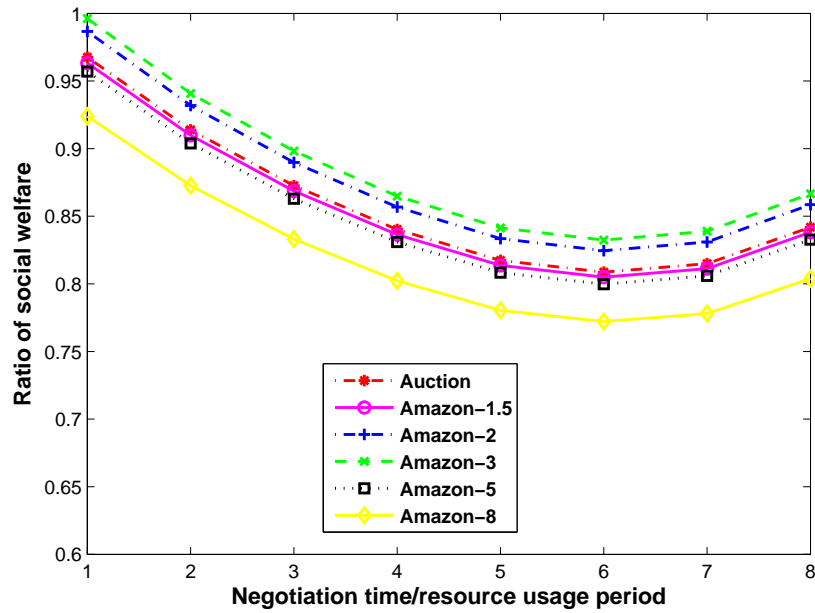


Figure 6.10. Social welfare and negotiation time

percentage (e.g., $\{0.1, 0.3, 0.5\}$) of a contract price. Experimental results show that *CRA* with decommitment is better than *CRA* without decommitment by no more than 3.5%.

6.6.3.2 Evaluating agents' negotiation strategies

Observation 6: Since it is impossible to find out agents' equilibrium strategies in the complex bargaining game, we designed strategies for agents by taking into account some important factors which are considered in the literature. While negotiation agents with the strategies achieved higher social welfare than other mechanisms, one may ask whether agents have an incentive to switch to other strategies. To answer this question, we tried some other strategies as follows: 1) each buyer makes only one set of agreements, 2) each buyer makes three sets of agreements, 3) when an agent decide to decommit from an agreement o , it decommits before $est(o)$ rather than decommits immediately, 4) and a seller makes contracts beyond its capability.

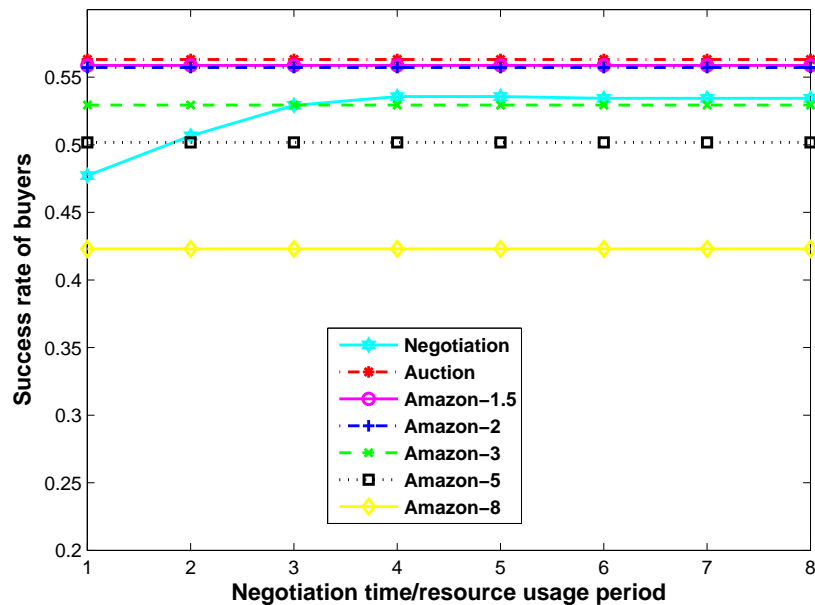


Figure 6.11. Success rate and negotiation time

We found that making these changes did not improve either utilities of agents with new strategies or *NG*'s performance. Always making two sets of contracts is a good choice due to the tradeoff between failing to finish the task and paying too much. While delaying decommitment does not “hurt” an agent directly, it hurts the agent “indirectly” since resource competition will increase if each agent holds more contracts. Further, it is better for a seller not to make contracts beyond its capability: if the resource competition is low, generally a seller cannot make contracts beyond its capability, and if the resource competition is high, a buyer is less likely to decommit from a contract and a seller may have to pay more penalties for contracts beyond its capability.

6.6.3.3 Sensitivity analysis

We also did additional experiments to explore the sensitivity of our experimental results to changes to the parameters of our experimental environments or assumptions about our negotiation model.

Observation 7: This work assumes that each agent knows the demand/supply ratio of each resource. In reality, an agent may not know the demand/supply ratio. We tested the negotiation model without this assumption and alternatively, each agent predicts the demand/supply ratio through its interaction with buyers. Specifically, a seller can estimate the competition of a resource according to 1) the requests for the resource from all the buyers in the last λ time points and 2) the total number of resources provided by other sellers. A buyer can estimate the competition of a resource according to bids from sellers. In this case, we found that the social welfare of *NG* is still 10% higher than other mechanisms.

Observation 8: This work also assumes that each agent knows each seller's cost of a resource. We found that the accuracy of this information does have a slight effect on agents' negotiation performance. When the believed cost is less than half of the actual cost, the average social welfare of *NG* is 6% lower than that of *NG* in which each buyer knows the actual cost.

Observation 9: Different from existing work on automated negotiation with recommitment, in our framework agents negotiate over both price and decommitment penalty. We compared setting penalties through negotiation with exogenous mechanisms for setting penalties [11, 98], e.g., fixed penalties (e.g., $\{0, 10, 20, 40\}$) or penalty as a percentage (e.g., $\{0.1, 0.3, 0.5\}$) of a contract price. We found that setting penalties through negotiation achieved much higher social welfare than those exogenous mechanisms for setting penalties. In fact, setting a decommitment penalty (function) to maximize social welfare is a difficult problem for the system designer due to lack of knowledge and agents' strategic behavior. Accordingly, it may be a good idea to give agents the flexibility to decide the decommitment penalty.

6.7 Summary

This chapter presents the design and implementation of a negotiation mechanism for dynamic resource allocation problem in cloud computing. Our negotiation model goes beyond the state of the art in the following aspects: 1) Multiple buyers and sellers are allowed to negotiate with each other concurrently and an agent is allowed to decommit from an agreement at the cost of paying a penalty; 2) Agents are allowed to negotiate over both price and penalty; 3) Negotiation strategies for both buyers and sellers consider important factors widely studied in the literature. Experimental results show that the proposed negotiation model outperforms different combinatorial auction mechanisms and Amazon's fixed price model. In general, the proposed mechanism can be applied in wide range of dynamic resource allocation problems. This chapter also complements previous work on leveled-commitment contracting by integrating a strategic contracting game with the leveled decommitting game and analyzing agents' equilibrium strategies in the contracting game.

CHAPTER 7

CONCLUSIONS AND FUTURE RESEARCH

This thesis set out to investigate the role of automated negotiation in various aspects of complex multi-agent resource allocation problems. In this final chapter, we will summarize the research contributions of this work, as well as discuss directions for future research.

7.1 Contributions

The work described in this thesis makes a number of important contributions to the state of the art in the area of agent mediated negotiation by extending both theoretical and heuristic bargaining approaches to more realistic settings involving uncertainty, market competition, decommitment, and acquirement of multiple resources. The contributions of this work can be summarized as follows:

- We present a novel algorithm to find pure strategy sequential equilibria in bilateral bargaining with multi-type uncertainty [6, 8]. Our algorithm combines together game theoretic analysis with search techniques. Our algorithm goes beyond existing algorithms dealing with complete information settings. Sequential games of incomplete information are ubiquitous and our approach is not specific to an application since it can be applied to other uncertainty settings, e.g., multi-issue negotiation with uncertain weight functions [51]. Our study shows that there exists at least one sequential equilibrium in more than 99.7% of scenarios we have tried in which there are deadline constraints and incomplete information. We also compared the performance of the equilibrium

strategies and representative heuristic based strategies. Empirical results show that agents with equilibrium strategies achieved higher utilities than agents with heuristic based strategies. Furthermore, when both agents adopt the equilibrium strategies, the agents achieved much higher social welfare than that in all other strategy combinations.

- We extend the alternating-offers protocol to handle multiple trading opportunities and market competition [7]. We provide an algorithm based on backward induction to compute the subgame perfect equilibrium of concurrent one-to-many negotiation and many-to-many negotiation. This is the first work on analyzing agents' equilibrium strategies in concurrent negotiation in markets. We observe that agents' bargaining power are affected by the proposing ordering and market competition. We find that for a large subset of the space of the parameters, agents' equilibrium strategies depend on the values of a narrow number of parameters. We also provide an algorithm to find a pure strategy sequential equilibrium in one-to-many negotiation where there is uncertainty regarding the reserve price of one agent.
- We develop and experimentally evaluate negotiation agents that negotiate for multiple resources where agents don't know the reserve price of each resource and are allowed to decommit from existing agreements. Existing work only considers single resource negotiation and often make unrealistic assumptions about agents' knowledge. The distinguishing feature of negotiation agents is that they are designed with the flexibility to adjust 1) the number of tentative agreements for each resource and 2) the amount of concession by reacting to i) changing market conditions, and ii) the current negotiation status of all concurrently negotiating threads. In addition, to avoid the risk of the "collapse" of the overall negotiation due to failing to acquire some scarce resources, nego-

tiation agents have the flexibility to adjust the deadline for different resources based on market competition, which allows agents to respond to uncertainties in resource planning. An extensive set of experiments were carried out and the experiments results show that each of the proposed heuristics contributes to improve agents' performance and our proposed approach achieved better negotiation results than representative samples of existing negotiation strategies.

- We propose a distributed negotiation mechanism for the problem of allocating networked resources in dynamic environment, such as cloud computing platforms. In the negotiation model, multiple buyers and sellers are allowed to negotiate with each other concurrently and an agent is allowed to decommit from an agreement at the cost of paying a penalty. Furthermore, agents negotiate over both a contract price and a decommitment penalty. We also propose negotiation strategies for both buyers and sellers considering important factors widely studied in the literature. Experimental results show the advantage of the negotiation model over different combinatorial auction mechanisms and Amazon's fixed price model. In general, the proposed mechanism can be applied in wide range of dynamic resource allocation problems. This is the first work that shows the importance of negotiation over decommitment penalties.

7.2 Future Research

Future research in negotiation theory includes attacking some challenging open negotiation problems concerning *issue multiplicity*, *information incompleteness*, and *negotiation involving multiple agents*. The problem of efficient negotiation over multiple issues is more difficult than simple issue negotiation due to difficulties in finding efficient mechanisms that produce Pareto optimal agreements. In presence of incomplete information, it is often difficult to compute agents' (sequential) equilibrium strategies. We have proposed a general algorithm dealing with bargaining with un-

certainty and there are several natural directions suggested by our research. The first one concerns the extension of our results in bargaining situations where there are two-sided uncertainty or other parameters are uncertain (e.g., discount factor). Second, we have seen from experimental results that there are more than one pure strategy sequential equilibrium in some scenarios. It would be useful to design coordination mechanisms for choosing certain equilibrium strategies for agents to play. In addition, characterizing bargaining games with no sequential equilibrium is also on the agenda. Finally, our experiments about the performance of equilibrium strategies thus far have focused on scenarios ranging from low to moderate complexity, but we wish to investigate much larger problems where there are longer deadlines and more buyer types. Regarding multiple agents, a central research topic in bargaining theory is understanding bargaining power, which is related to the relative abilities of agents in a situation to exert influence over each other. In bilateral bargaining, each agent's bargaining power is affected by its reserve price, patience attitude, deadline, etc. When many buyers and sellers are involved in negotiation, it is important to investigate how the market competition will affect agents' equilibrium bargaining strategies. With a large number of buyers and sellers, a single agent is unlikely to have much influence on the market equilibrium.

Another future research direction is looking at new applications of automated negotiation for complex multi-agent resource allocation. In practical multi-agent resource allocation problems (e.g., the two resource allocation problem discussed in Chapter 1), information incompleteness and existence of market competition make it intractable to compute agents' equilibrium strategies. The community has explored solutions which are alternative to the classic game theoretic solution by bounding agents' rationality. In such situations, designing heuristics that perform well is still a challenging problem. First, an agent needs to learn knowledge from its negotiation history. Second, market dynamics may require an agent to reason about future

trading opportunities. In addition, each agent needs to reason about other agents' strategies. I will try to apply the distributed negotiation approach for resource allocation in new applications I have not looked into before, e.g., electronic commerce, supply chain, web/grid service composition, workflow, and enterprise integration.

In addition to designing negotiation strategies to maximizing an agent's utility, designing negotiation mechanisms that maximize some global performance measures (e.g., social welfare) is also a future research direction. One line of research is investigating some simplified bargaining games. For instance, Athey and Segal [16] consider the bargaining mechanism design problem of allocating a good between two players where the players' valuations in each period are private information, and the valuations change over time following a first-order Markov process. The other line of research is considering more complex environment and evaluating different mechanisms through experimentation. The negotiation model in Chapter 6 falls into this category and there are a number of future research directions. First, in the current design, an agent will make its decision (e.g., accept, confirm) immediately after it receives a message. As future work, we will consider the role of delaying making decisions. Second, while it is impossible to derive agents' equilibrium strategies in such dynamic resource allocation game, it would be interesting to investigate agents' rational strategies in some simplified scenarios [7]. Third, fully understanding the role of decommitment in this resource allocation game deserves further analysis and experimentation. Finally, studying and evaluating other auction models (e.g., partially centralized auction models with different ways of determining when to run auctions) are also necessary.

Another interesting future research direction is bargaining in trading networks. Different from trading in markets, a buyer and a seller can negotiate for an agreement if and only if they have a relationship, or "link", to engage in exchange. This setting is practical since individual buyers and sellers often trade through intermediaries,

not all buyers and sellers have access to the same intermediaries. One good example of this setting is the trade of agricultural goods in developing countries [25]. Given inadequate transportation networks, and poor farmers' limited access to capital, many farmers have no alternative to trading with middlemen in inefficient local markets. Bargaining in trading networks has received a lot of attention in recent years (e.g., [25, 103, 76, 32, 38]) and the focus is analyzing agents' strategic behavior. Future research should include analyzing agents' strategies in incomplete information settings and how network structure will agents' strategies and bargaining outcome.

Bargaining theory so far assumes that agents are fully rational, which rarely holds in real-world domains as is well known, such human beings may not be utility maximizers. Therefore, theoretic analysis may be not useful in practice. Future research will include analyzing strategic behavior of agents with bounded rationality. Another related future research direction is building systems to support human negotiation (e.g., [88, 89]), which is difficult due to a number of reasons. First, we need to consider much larger negotiation space and strategy space. For instance, human beings often use body languages while doing negotiation. Second, we need to consider many other factors such as emotion, trust, power, culture, belief, desire, and intention.

Future research should also include designing new business models for creating virtual enterprises. A virtual enterprise is a temporary group of fully autonomous agents that is formed to meet a special objective or to provide a special service. Achieving this objective or service involves performing a series of tasks that require repeated interactions among virtual enterprise members.

BIBLIOGRAPHY

- [1] GENI System Overview. <http://groups.geni.net/geni/wiki/GeniSysOvrvw>.
- [2] NSF Global Environment For Network Innovations (GENI). <http://www.geni.net/>.
- [3] Amazon Elastic Compute Cloud. <http://aws.amazon.com/ec2/>.
- [4] Aknine, Samir, Pinson, Suzanne, and Shakun, Melvin. An extended multi-agent negotiation protocol. *Autonomous Agents and Multi-Agent Systems* 8, 1 (2004), 5–45.
- [5] An, Bo, Douglis, Fred, and Ye, Fan. Heuristics for negotiation schedules in multi-plan optimization. In *Proc. of the Seventh International Joint Conference on Autonomous Agents and Multi-Agent Systems* (2008), pp. 551–558.
- [6] An, Bo, Gatti, Nicola, and Lesser, Victor. Bilateral bargaining with one-sided two-type uncertainty. In *Proc. of the 9th IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT)* (Sep. 2009), pp. 403–410.
- [7] An, Bo, Gatti, Nicola, and Lesser, Victor. Extending alternating-offers bargaining in one-to-many and many-to-many settings. In *Proc. of the 9th IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT)* (Sep. 2009), pp. 423–426.
- [8] An, Bo, Gatti, Nicola, and Lesser, Victor. Searching for pure strategy equilibria in bilateral bargaining with one-sided uncertainty. In *Proc. of the 9th International Joint Conference on Autonomous Agents and Multi-Agent Systems* (May 2010), pp. 1607–1608.
- [9] An, Bo, and Lesser, Victor. Characterizing contract-based multi-agent resource allocation in networks. *IEEE Trans. on Systems, Man and Cybernetics, Part B: Cybernetics* 40, 3 (June 2010), 575–586.
- [10] An, Bo, Lesser, Victor, Irwin, David, and Zink, Michael. Automated negotiation with decommitment for dynamic resource allocation in cloud computing. *Proc. of the Ninth International Joint Conference on Autonomous Agents and Multi-Agent Systems* (May 2010), 981–988.
- [11] An, Bo, Lesser, Victor, and Sim, Kwang Mong. Decommitment in multi-resource negotiation. In *Proc. of the International Joint Conference on Autonomous Agents and Multi-Agent Systems* (2008), pp. 1553–1556.

- [12] An, Bo, Lesser, Victor, and Sim, Kwang Mong. Strategic agents for multi-resource negotiation. *Journal of Autonomous Agents and Multi-Agent Systems* (to appear).
- [13] An, Bo, Sim, Kwang Mong, and Lesser, Victor. Evolving the best-response strategy to decide when to make a proposal. In *Proc. of the 2007 IEEE Congress on Evolutionary Computation* (2007), pp. 1035–1042.
- [14] An, Bo, Sim, Kwang Mong, Tang, Lianggui, Li, Shuangqing, and Cheng, Daijie. A continuous time negotiation mechanism for software agents. *IEEE Trans. on Systems, Man and Cybernetics, Part B* 36, 6 (Dec. 2006), 1261–1272.
- [15] Andersson, Martin, and Sandholm, Tuomas. Leveled commitment contracts with myopic and strategic agents. *Journal of Economic Dynamics & Control* 25 (2001), 615–640.
- [16] Athey, Susan, and Segal, Ilya. Designing efficient mechanisms for dynamic bilateral trading games. *American Economic Review* 97, 2 (2007), 131–136.
- [17] Aumann, R. *Frontiers of Economics*. Basil Blackwell, Oxford, 1985, ch. What is game theory trying to accomplish?, pp. 909–924.
- [18] Ausubel, L. M., Cramton, P. C., and Deneckere, R. J. *Handbook of Game Theory*, vol. 3. Elsevier Science, 2002, ch. Bargaining with incomplete information, pp. 1897–1945.
- [19] Auyoung, Alvin Alvin, Chun, Brent N., Snoeren, Alex C., and Vahdat, Amin. Resource allocation in federated distributed computing infrastructures. In *Proc. of the 1st Workshop on Operating System and Architectural Support for the On-demand IT Infrastructure (OASIS 2004)* (2004).
- [20] Bac, M., and Raff, H. Issue-by-issue negotiations: the role of information and time preference. *Games and Economic Behavior* 13 (1996), 125–134.
- [21] Bergemann, Dirk, and Said, Maher. Dynamic auctions: a survey. Cowles Foundation Discussion Papers 1757, Cowles Foundation, Yale University, Mar. 2010.
- [22] Bernstein, Daniel. *Complexity Analysis and Optimal Algorithms for Decentralized Decision Making*. PhD thesis, University of Massachusetts, Amherst, 2005.
- [23] Bikhchandani, Sushil. A bargaining model with incomplete information. *Review of Economic Studies* 59, 1 (January 1992), 187–203.
- [24] Binmore, K., Shaked, A., and Sutton, J. An outside option experiment. *The Quarterly Journal of Economics* 104, 4 (1989), 753–770.
- [25] Blume, Larry, Easley, David, Kleinberg, Jon, and Tardos, Eva. Trading networks with price-setting agents. In *Proc. of the 8th ACM conference on Electronic commerce* (2007), pp. 143–151.

- [26] Blumrosen, Liad, and Nisan, Noam. *Algorithmic Game Theory*. Cambridge University Press, 2007, ch. Combinatorial Auctions.
- [27] Branson, Michael, Douglass, Fred, Fawcett, Bradley, Liu, Zhen, Riabov, Anton, and Ye, Fan. CLASP: collaborating, autonomous data stream processing systems. In *Proc. of the ACM/IFIP/USENIX 8th International Middleware Conference (2007)*, pp. 348–367.
- [28] Brazier, F., Cornelissen, F., Gustavsson, R., Jonker, C., Lindeberg, O., Polak, B., and Treur, J. A multi-agent system performing one-to-many negotiation for load balancing of electricity use. *Electronic Commerce Research and Applications 1, 2* (2002), 208–224.
- [29] Busch, L., and Horstmann, I. Bargaining frictions, bargaining procedures and implied costs in multiple-issue bargaining. *Economica 64* (1997), 669–680.
- [30] Busch, L., and Horstmann, I. Endogenous incomplete contracts: a bargaining approach. *Games and Economic Behavior 19* (1997), 144–148.
- [31] Busch, L., and Horstmann, I. Signaling via an agenda in multi-issue bargaining with incomplete information. *Economic Theory 13* (1999), 561–575.
- [32] Charness, Gary, Corominas-Bosch, Margarida, and Frechette, Guillaume R. Bargaining and network structure: an experiment. *Journal of Economic Theory 136, 1* (2007), 28–65.
- [33] Chatterjee, K., and Lee, Ching Chyi. Bargaining and search with incomplete information about outside options. *Games and Economic Behavior 22* (1998), 203–237.
- [34] Chatterjee, K., and Samuelson, L. Bargaining with two-sided incomplete information: an infinite horizon model with alternating offers. *Review of Economic Studies 54, 2* (1987), 175–192.
- [35] Chatterjee, K., and Samuelson, L. Bargaining under two-sided incomplete information: the unrestricted offers case. *Operations Research 36* (1988), 605–618.
- [36] Chatterjee, K., and Samuelson, W. Bargaining under incomplete information. *Operations Research 31* (1983), 835–851.
- [37] Coehoorn, Robert M., and Jennings, Nicholas R. Learning an opponent’s preferences to make effective multi-issue negotiation trade-offs. In *Proc. 6th International Conference on Electronic Commerce (2004)*, pp. 59–68.
- [38] Corominas-Bosch, Margarida. Bargaining in a network of buyers and sellers. *Journal of Economic Theory 115, 1* (2004), 35–77.
- [39] Cramton, P. Bargaining with incomplete information: an infinite-horizon model with two-sided uncertainty. *Review of Economic Studies 51, 4* (1984), 579–593.

- [40] Cramton, P. *Game Theoretic Models of Bargaining*. Cambridge University Press, 1985, ch. Sequential bargaining mechanisms, pp. 149–179.
- [41] Cramton, P. Strategic delay in bargaining with two-sided uncertainty. *Review of Economic Studies* 59, 1 (1992), 205–225.
- [42] Cramton, Peter, Shoham, Yoav, and Steinberg, Richard, Eds. *Combinatorial Auctions*. MIT Press, 2006.
- [43] Dang, Jiangbo, and Huhns, Michael. An extended protocol for multiple-issue concurrent negotiation. In *Proc. of the 20th national conference on Artificial intelligence* (2005), pp. 65–70.
- [44] Dantzig, George. Discrete variable extremum problems. *Operations Research* 5 (1957), 266–277.
- [45] Dunne, P. E. Extremal behaviour in multiagent contract negotiation. *Journal of Artificial Intelligence Research* 23 (2005), 41–78.
- [46] Ehtamo, Harri, Kettunen, Eero, and Hamalainen, Raimo P. Searching for joint gains in multi-party negotiations. *European Journal of Operational Research* 130, 1 (April 2001), 54–69.
- [47] Endriss, Ulle, and Maudet, Nicolas. On the communication complexity of multi-lateral trading: extended report. *Journal of Autonomous Agents and Multiagent Systems* 11, 1 (2005), 91–107.
- [48] Faratin, P., Sierra, C., and Jennings, N. R. Negotiation decision functions for autonomous agents. *Int. Journal of Robotics and Autonomous Systems* 24, 3-4 (1998), 159–182.
- [49] Faratin, P., Sierra, C., and Jennings, N. R. Using similarity criteria to make issue trade-offs in automated negotiations. *Artificial Intelligence* 142, 2 (2002), 205–237.
- [50] Fatima, Shaheen S., Wooldridge, Michael, and Jennings, Nicholas R. An agenda-based framework for multi-issue negotiation. *Artificial Intelligence* 152, 1 (2004), 1–45.
- [51] Fatima, Shaheen S., Wooldridge, Michael, and Jennings, Nicholas R. Multi-issue negotiation with deadlines. *Journal of Artificial Intelligence Research* 27 (2006), 381–417.
- [52] Fatima, Shaheen S., Wooldridge, Michael, and Jennings, Nicholas R. On efficient procedures for multi-issue negotiation. In *Proc. of the Trading Agent Design and Analysis and Agent Mediated Electronic Commerce* (2006), pp. 71–84.

- [53] Fershtman, C. The importance of the agenda in bargaining. *Games and Economic Behavior* 2 (1990), 224–238.
- [54] Flamini, Francesca. First things first? the agenda formation problem for multi-issue committees. *Journal of Economic Behavior & Organization* 63, 1 (2007), 138–157.
- [55] Fu, Yun, Chase, Jeffrey, Chun, Brent, Schwab, Stephen, and Vahdat, Amin. Sharp: an architecture for secure resource peering. In *Proc. of the 19th ACM Symposium on Operating System Principles* (2003), pp. 133–148.
- [56] Fudenberg, D., Mobius, Markus, and Szeidl, Adam. Existence of equilibrium in large double auctions. *Journal of Economic Theory* 127, 1 (2007), 550–567.
- [57] Fudenberg, D., and Tirole, J. Sequential bargaining under incomplete information. *Review of Economic Studies* 50, 2 (1983), 221–248.
- [58] Fudenberg, D., and Tirole, J. *Game Theory*. MIT Press, Cambridge, MA, 1991.
- [59] Fudenberg, Drew, Levine, David K., and Tirole, Jean. Incomplete information bargaining with outside opportunities. *The Quarterly Journal of Economics* 102, 1 (1987), 37–50.
- [60] Gantner, Anita. Bargaining, search, and outside options. *Games and Economic Behavior* 62, 2 (2008), 417–435.
- [61] Gatti, Nicola, Giunta, Francesco Di, and Marino, Stefano. Alternating-offers bargaining with one-sided uncertain deadlines: an efficient algorithm. *Artificial Intelligence* 172, 8-9 (2008), 1119–1157.
- [62] Gilpin, Andrew, and Sandholm, Tuomas. Finding equilibria in large sequential games of imperfect information. In *Proceedings of the 7th ACM conference on Electronic commerce* (2006), pp. 160–169.
- [63] Gilpin, Andrew, and Sandholm, Tuomas. Lossless abstraction of imperfect information games. *Journal ACM* 54, 5 (2007).
- [64] Gimpel, H., Jennings, N. R., Kersten, G., Okenfels, A., and Weinhardt, C. *Negotiation, Auctions, and Market Engineering*. Springer, 2008, ch. Market engineering: a research agenda, pp. 1–15.
- [65] Giunta, Francesco, and Gatti, Nicola. Bargaining over multiple issues in finite horizon alternating-offers protocol. *Annals of Mathematics and Artificial Intelligence* 47, 3-4 (2006), 251–271.
- [66] In, Y., and Serrano, R. Agenda restrictions in multi-issue bargaining. *Journal of Economic Behavior and Organization* 53 (2004), 385–399.
- [67] Inderst, R. Multi-issue bargaining with endogenous agenda. *Games and Economic Behavior* 30 (2000), 64–82.

- [68] Irwin, David, Chase, Jeffrey, Grit, Laura, Yumerefendi, Aydan, and Becker, David. Sharing networked resources with brokered leases. In *Proc. of the 2006 USENIX Annual Technical Conference (2006)*, pp. 199–212.
- [69] Jain, Navendu, Amini, Lisa, Andrade, Henrique, King, Richard, Park, Yoonho, Selo, Philippe, and Venkatramani, Chitra. Design, implementation, and evaluation of the linear road benchmark on the stream processing core. In *Proc. of the 25th ACM SIGMOD International Conference on Management of Data (2006)*, pp. 431–442.
- [70] Jennings, N. R., Faratin, P., Lomuscio, A. R., Parsons, S., Sierra, C., and Wooldridge, M. Automated negotiation: prospects, methods and challenges. *Group Decision and Negotiation* 10, 2 (2001), 199–215.
- [71] Jennings, N. R., Norman, T. J., Faratin, P., O’Brien, P., and Odgers, B. Autonomous agents for business process management. *Applied Artificial Intelligence* 14, 2 (2000), 145–189.
- [72] Kersten, G., Noronha, S., and Teich, J. Are all e-commerce negotiations auctions? In *Proc. of the 4th International Conference on the Design of Cooperative Systems (2000)*, pp. 535–541.
- [73] Klein, Mark, Faratin, Peyman, Sayama, Hiroki, and Bar-yam, Yaneer. Negotiating complex contracts. *Group Decision and Negotiation* 12, 2 (2003), 111–125.
- [74] Koller, Daphne, Megiddo, Nimrod, and von Stengel, Bernhard. Fast algorithms for finding randomized strategies in game trees. In *Proceedings of the twenty-sixth Annual ACM Symposium on Theory of Computing (1994)*, pp. 750–759.
- [75] Krainin, Michael, An, Bo, and Lesser, Victor. An application of automated negotiation to distributed task allocation. In *Proceedings of the 2007 IEEE/WIC/ACM International Conference on Intelligent Agent Technology (Nov. 2007)*, pp. 138–145.
- [76] Kranton, Rachel E., and Minehart, Deborah F. A theory of buyer-seller networks. *American Economic Review* 91, 3 (2001), 485–508.
- [77] Kraus, S. *Strategic Negotiation in Multiagent Environments*. MIT Press, Cambridge, MA, 2001.
- [78] Kraus, S., J.Wilkenfeld, and Zlotkin, G. Multiagent negotiation under time constraints. *Artificial Intelligence* 75, 2 (1995), 297–345.
- [79] Kreps, David, and Wilson, Robert. Sequential equilibria. *Econometrica* 50, 4 (1982), 863–894.
- [80] Lai, Guoming, Li, Cuihong, and Sycara, Katia. Efficient multi-attribute negotiation with incomplete information. *Group Decision and Negotiation* 15 (2006), 511–528.

- [81] Lai, Guoming, Li, Cuihong, Sycara, Katia, and Giampapa, Joseph. Literature review on multi-attribute negotiations. Tech. Rep. CMU-RI-TR-04-66, Robotics Institute, Carnegie Mellon University, December 2004.
- [82] Lai, Guoming, Sycara, Katia, and Li, Cuihong. A decentralized model for automated multi-attribute negotiations with incomplete information and general utility functions. *Multiagent and Grid Systems* 4, 1 (2008), 45–65.
- [83] Lai, K., Rasmusson, L., E. Adar, L. Zhang, and Huberman, B. Tycoon: An implementation of a distributed, market-based resource allocation system. *Multiagent and Grid Systems* 1, 3 (2005), 169–182.
- [84] Lang, K., and Rosenthal, R. W. Bargaining piecemeal or all at once. *Multiagent and Grid Systems* 111 (2001), 526–540.
- [85] Lau, R. Y.K. Towards genetically optimised multi-agent multi-issue negotiations. In *Proc. of the 38th Hawaii Int. Conference on System Sciences* (Jan. 2005), IEEE Press.
- [86] Lee, Cynthia, and Snavely, Allan. Precise and realistic utility functions for user-centric performance analysis of schedulers. In *Proc. of the 16th Int. Symp. on High Performance Distributed Computing* (2007), pp. 107–116.
- [87] Li, Cuihong, Giampapa, Joseph Andrew, and Sycara, Katia. Bilateral negotiation decisions with uncertain dynamic outside options. *IEEE Trans. on Systems, Man, and Cybernetics, Part C: Special issue on Game-theoretic Analysis and Stochastic Simulation of Negotiation Agents* 36, 1 (Jan. 2006), 31–44.
- [88] Lin, Raz, and Kraus, Sarit. Can automated agents proficiently negotiate with humans? *Communications of the ACM* 53, 1 (2010), 78–88.
- [89] Lin, Raz, Oshrat, Yinon, and Kraus, Sarit. Investigating the benefits of automated negotiations in enhancing people’s negotiation skills. In *Proc. of the 8th International Conference on Autonomous Agents and Multiagent Systems* (2009), pp. 345–352.
- [90] Lomuscio, A. R., Wooldridge, M., and Jennings, N. R. A classification scheme for negotiation in electronic commerce. *Int. Journal Group Decision Negotiation* 12, 1 (Jul./Aug. 2003), 31–56.
- [91] Matos, N., Sierra, C., and Jennings, N. R. Determining successful negotiation strategies: an evolutionary approach. In *Proc. of the 3rd Int. Conference on Multi-Agent Systems* (Jan. 1998), pp. 182–189.
- [92] Miltersen, P., and Sorensen, T. Computing sequential equilibria for two-player games. In *Proc. of the ACMCSIAM Symposium on Discrete Algorithm (SODA)* (2006), pp. 107–116.

- [93] Muthoo, Abhinay. On the strategic role of outside options in bilateral bargaining. *Operations Research* 43, 2 (1995), 292–297.
- [94] Myerson, R., and Satterthwaite, M. Efficient mechanism for bilateral trading. *Journal of Economic Theory* 28 (1983), 265–281.
- [95] Nash, J. The bargaining problem. *Econometrica* 18, 2 (1950), 155–162.
- [96] Nguyen, T. D., and Jennings, N. R. A heuristic model of concurrent bi-lateral negotiations in incomplete information settings. In *Proc. of the eighteenth Int. Joint Conference on Artificial Intelligence* (2003), pp. 1467–1469.
- [97] Nguyen, T. D., and Jennings, N. R. Coordinating multiple concurrent negotiations. In *Proc. of the 3rd International Conference on Autonomous Agents and Multi-Agent Systems* (2004), pp. 1064–1071.
- [98] Nguyen, T. D., and Jennings, N. R. Managing commitments in multiple concurrent negotiations. *Electronic Commerce Research and Applications* 4, 4 (2005), 362–376.
- [99] Nisan, Noam, Roughgarden, Tim, Tardos, Eva, and Vazirani, Vijay V., Eds. *Algorithmic Game Theory*. Cambridge Univ. Press, 2007.
- [100] Osborne, M. J., and Rubinstein, A. *Bargaining and markets*. Academic, New York, 1990.
- [101] Parkes, David. *Algorithmic Game Theory*. Cambridge University Press, 2007, ch. Online Mechanisms.
- [102] Perry, Motty. An example of price formation in bilateral situations: a bargaining model with incomplete information. *Econometrica* 54, 2 (1986), 313–321.
- [103] Polanski, Arnold. Bilateral bargaining in networks. *Journal of Economic Theory* 134, 1 (2007), 557–565.
- [104] Porter, Ryan, Nudelman, Eugene, and Shoham, Yoav. Simple search methods for finding a nash equilibrium. *Games and Economic Behavior* 63 (2008), 642–662.
- [105] Preist, Chris, Bartolini, Claudio, and Hyde, Andrew. Agent-based service composition through simultaneous negotiation in forward and reverse auctions. In *Proc. of the 4th ACM Conference on Electronic Commerce* (2003), pp. 55–63.
- [106] Rahwan, I., Kowalczyk, R., and Pham, H. H. Intelligent agents for automated one-to-many e-commerce negotiation. *Australian Computer Science Communications* 24, 1 (Jan./Feb. 2002), 197–204.
- [107] Raiffa, H. *The art and science of negotiation*. Harvard Univ. Press, Cambridge, MA, 1982.

- [108] Riabov, Anton, and Liu, Zhen. Planning for stream processing systems. In *Proc. of the Twentieth National Conference on Artificial Intelligence* (2005), pp. 1205–1210.
- [109] Rosenschein, J., and Zlotkin, G. *Rules of Encounter: Designing conventions for automated negotiation among computers*. MIT Press, Cambridge, MA, 1994.
- [110] Rothkopf, Michael, Pekec, Aleksandar, and Harstad, Ronald. Computationally manageable combinatorial auctions. *Management Science* 44, 8 (1998), 1131–1147.
- [111] Rubinstein, A. Perfect equilibrium in a bargaining model. *Econometrica* 50, 1 (1982), 97–109.
- [112] Rubinstein, A. A bargaining model under incomplete information about time preferences. *Econometrica* 53, 5 (1985), 1151–1172.
- [113] Sandholm, T. Contract types for satisficing task allocation: I theoretical results. In *Proc. of the AAAI Spring Symposium* (1998).
- [114] Sandholm, T. Expressive commerce and its application to sourcing: how we conducted \$35 billion of generalized combinatorial auctions. *AI Magazine* 28, 3 (2007), 45–58.
- [115] Sandholm, T., and Lesser, V. Issues in automated negotiation and electronic commerce: Extending the contract net framework. In *First International Conference on Multi-Agent Systems*, (1995), pp. 328–335.
- [116] Sandholm, T., Sikka, S., and Norden, S. Algorithms for optimizing leveled commitment contracts. In *Proc. of the 16th International Joint Conference on Artificial Intelligence (IJCAI)* (1999), pp. 535–541.
- [117] Sandholm, T., and Vulkan, N. Bargaining with deadlines. In *Proc. of the 16th National Conference on Artificial Intelligence (AAAI)* (1999), pp. 44–51.
- [118] Sandholm, T. W. An implementation of the contract net protocol based on marginal cost calculations. In *Proc. of the 12th International Workshop on Distributed Artificial Intelligence* (1993), pp. 295–308.
- [119] Sandholm, Tuomas. Automated negotiation. *Communication of the ACM* 42, 3 (1999), 84–85.
- [120] Sandholm, Tuomas. Agents in electronic commerce: component technologies for automated negotiation and coalition formation. *Autonomous Agents and Multi-Agent Systems* 3, 1 (2000), 73–96.
- [121] Sandholm, Tuomas, and Lesser, Victor. Leveled commitment contracts and strategic breach. *Games and Economic Behavior* 35 (2001), 212–270.

- [122] Sandholm, Tuomas, and Zhou, Y. Surplus equivalence of leveled commitment contracts. *Artificial Intelligence* 142 (2002), 239–264.
- [123] Shaked, Avner, and Sutton, John. Involuntary unemployment as a perfect equilibrium in a bargaining model. *Econometrica* 52, 6 (November 1984), 1351–1364.
- [124] Sim, K. M. A market-driven model for designing negotiation agents. *Computational Intelligence, Special issue in Agent Technology for E-commerce* 18, 4 (2002), 618–637.
- [125] Sim, K. M. Equilibria, prudent compromises, and the “waiting” game. *IEEE Trans. on Systems, Man and Cybernetics, Part B* 33, 3 (Jun. 2005), 712–724.
- [126] Sim, K. M. G-commerce, market-driven g-negotiation agents and grid resource management. *IEEE Transactions on Systems, Man and Cybernetics, Part B* 36, 6 (2006), 1381–1394.
- [127] Sim, K. M., and Choi, C. Y. Agents that react to changing market situations. *IEEE Trans. on Systems, Man and Cybernetics, Part B* 33, 2 (Apr. 2003), 188–201.
- [128] Sim, K. M., and Shi, B. Concurrent negotiation and coordination for grid resource co-allocation. *IEEE Trans. on Systems, Man and Cybernetics, Part B* 40, 2 (2009).
- [129] Sim, K. M., and Wang, S.Y. Flexible negotiation agent with relaxed decision rules. *IEEE Trans. on Systems, Man and Cybernetics, Part B* 34, 3 (Jun. 2004), 1602–1608.
- [130] Sim, Kwang Mong, and Shi, Benyun. Adaptive commitment management strategy profiles for concurrent negotiations. In *Proc. of The first International Workshop on Agent-based Complex Automated Negotiations* (2008), pp. 16–23.
- [131] Smith, R. G. The contract net protocol: high-level communication and control in a distributed problem solver. *IEEE Transactions on Computers* C-29, 12 (August 1980), 1104–1113.
- [132] Stahl, I. *Bargaining Theory*. Stockholm School of Economics, Stockholm, Sweden, 1972.
- [133] Sycara, K. Problem restructuring in negotiation. *Management Science* 37, 10 (1991), 1248–1268.
- [134] Vickrey, William. Counterspeculation and competitive sealed tenders. *Journal of Finance* 16, 1 (1961), 8–37.
- [135] von Stengel, Bernhard. *Algorithmic Game Theory*. Cambridge Univ. Press, 2007, ch. Equilibrium Computation for Two-Player Games in Strategic and Extensive Form.

- [136] Zeng, D., and Sycara, K. Bayesian learning in negotiation. *Int. Journal of Human Computer Systems* 48, 3-4 (1998), 125–141.
- [137] Zhang, Xiaoqin, Lesser, Victor, and Abdallah, Sherief. Efficient management of multi-linked negotiation based on a formalized model. *Journal of Autonomous Agents and MultiAgent Systems* 10, 2 (2005), 165–205.
- [138] Zhang, Xiaoqin, Lesser, Victor, and Wagner, Tom. Integrative negotiation among agents situated in organizations. *IEEE Transactions on System, Man, and Cybernetics: Part C* 36, 1 (2006), 19–30.