


Interactions in Multiagent Systems: Fairness, Social Optimality and Individual Rationality

郝建业 梁浩锋

多智能代理系统中的交互
——公平性、社会最优与个体理性(英文版)



高等教育出版社

新一代信息科学与技术丛书

in Multiagent Systems: Fairness, Social Optimality and Individual Rationality

多智能代理系统中的交互
——公平性、社会最优与个体理性 (英文版)

郝建业 梁浩锋

高等教育出版社·北京

图书在版编目(CIP)数据

多智能代理系统中的交互：公平性、社会最优与个体理性 = Interactions in Multiagent Systems: Fairness, Social Optimality and Individual Rationality : 英文 / 郝建业, 梁浩锋著. -- 北京 : 高等教育出版社, 2016. 6

ISBN 978-7-04-044111-6

I. ①多… II. ①郝… ②梁… III. ①智能系统-研究-英文 IV. ①TP18

中国版本图书馆CIP数据核字(2016)第124025号

策划编辑 刘英
责任印制 刘思涵

责任编辑 刘英

封面设计 张志

出版发行 高等教育出版社
社址 北京市西城区德外大街4号
邮政编码 100120
印刷 山东鸿君杰文化发展有限公司
开本 787mm×1092mm 1/16
印张 12
字数 240千字
购书热线 010-58581118
咨询电话 400-810-0598

网 址 <http://www.hep.edu.cn>
<http://www.hep.com.cn>
网上订购 <http://www.hepmall.com.cn>
<http://www.hepmall.com>
<http://www.hepmall.cn>

版次 2016年6月第1版
印次 2016年6月第1次印刷
定价 59.00元

本书如有缺页、倒页、脱页等质量问题, 请到所购图书销售部门联系调换
版权所有 侵权必究
物料号 44111-00

本书由高等教育出版社和 Springer 公司合作出版, 由 Springer 公司负责编辑加工和排版, 故书中量和单位以及图、表等难免存在不符合我国编辑规范之处。特此说明。

Interactions in Multiagent Systems: Fairness, Social Optimality and Individual Rationality

Preface

It has long been a fundamental assumption in the research of autonomous agents and multiagent systems that intelligent agents should be rational—well, at least as rational as possible even if they are handicapped, like lacking access to all information that is needed for rational decision-making. While this assumption is only natural and right, rationality is actually only a better word to use if we try to avoid the use of the word selfishness to describe the deed of concerning only with one's own interests and being regardless of that of others.

When we first decided to work on the issues of fairness and social optimality in multiagent systems, with due respect to individual rationality of agents, we knew that we were facing a nontrivial challenge. The first difficulty that we encountered was the definition of fairness. We did spend some time to investigate this issue, delving into the definitions given by various authoritative dictionaries, hoping to obtain some mathematical formalism that captured the essence of these definitions. Unfortunately, the definitions in most of these dictionaries were not too helpful to us. While all our efforts seemed to have ended up in vain, we found that we had no choice but to start by taking a simplistic approach to reduce fairness to equality of agents' utility values (in some cases with a sufficiently small difference). Such an approach suffers from several drawbacks, as it is not able to capture many real-life situations. For example, in real life some people might consider it fair for rich people to receive less and poor people to receive more concerning the allocation of public resources. To this end, we looked into fairness models considering other types of fairness such as inequity-aversion and reciprocity. We hope the work presented in this book can shed light on this research area, which could attract more people to continue to work on it.

In this book we aim at tackling the issues of fairness and social optimality in multiagent systems from several aspects. We start with a relatively simple setting, in which agents use adaptive periodical strategy for coordination toward fairness. Agents interact with one another and review and adjust their strategies periodically. We then proceed to investigate how agents can coordinate to social optimality with reinforcement social learning in various types of games. After that, we discuss a related scenario of automated agent negotiation and present a

negotiation strategy that aims at achieving a mutually beneficial agreement. We also considered a competitive negotiation problem and present an adaptive negotiation strategy enabling a rational agent to maximize his benefits from negotiation. Finally, we propose the novel idea of decision entrustment in two-agent repeated games—that is, an agent is allowed the option of entrusting his opponent to make joint decisions in bilateral interaction—and show that social optimality sustained by Nash equilibrium can be achieved.

There are several assumptions that we employ in the work presented in this book. The first one is that agents can always learn from their experiences, and the learning is always rational. That is to say, if an agent, through his experience, finds that an action is most likely to bring its best utility under certain conditions, then the agent should in principle use that action when conditions are met. This is most appropriate when the environment is cooperative. Second, agents should be able to observe the other agents—or some of them—in the system, including the other agents' actions and even the utilities they obtain through interaction. Last but not least, we base our work on some principles proposed by cognitive psychologists. For example, we use a theory by Fehr and Schmidt that people are inequality averse under certain conditions, as well as the research result due to Dawes and Thaleri that people are willing to be kind to those who are being kind and harsh to those who are being unkind. These research works provide a solid cognitive psychological basis for many of the assumptions used throughout this book.

Finally, we would like to express our gratitude for all the help from Higher Education Press and Springer in the publication of this book.

Tianjin, China
Hong Kong, China
August 2015

Jianye Hao
Ho-fung Leung

郑重声明

高等教育出版社依法对本书享有专有出版权。任何未经许可的复制、销售行为均违反《中华人民共和国著作权法》，其行为人将承担相应的民事责任和行政责任；构成犯罪的，将被依法追究刑事责任。为了维护市场秩序，保护读者的合法权益，避免读者误用盗版书造成不良后果，我社将配合行政执法部门和司法机关对违法犯罪的单位和个人进行严厉打击。社会各界人士如发现上述侵权行为，希望及时举报，本社将奖励举报有功人员。

反盗版举报电话 (010) 58581999 58582371 58582488

反盗版举报传真 (010) 82086060

反盗版举报邮箱 dd@hep.com.cn

通信地址 北京市西城区德外大街4号 高等教育出版社法律事务与版权管理部

邮政编码 100120

Contents

1	Introduction	1
1.1	Overview of the Chapters	2
1.2	Guide to the Book	3
	References	5
2	Background and Previous Work	7
2.1	Background	7
2.1.1	Single-Shot Normal-Form Game	7
2.1.2	Repeated Games	9
2.2	Cooperative Multiagent Systems	10
2.2.1	Achieving Nash Equilibrium	10
2.2.2	Achieving Fairness	13
2.2.3	Achieving Social Optimality	16
2.3	Competitive Multiagent Systems	19
2.3.1	Achieving Nash Equilibrium	19
2.3.2	Maximizing Individual Benefits	20
2.3.3	Achieving Pareto-Optimality	21
	References	23
3	Fairness in Cooperative Multiagent Systems	27
3.1	An Adaptive Periodic Strategy for Achieving Fairness	28
3.1.1	Motivation	28
3.1.2	Problem Specification	30
3.1.3	An Adaptive Periodic Strategy	32
3.1.4	Properties of the Adaptive Strategy	36
3.1.5	Experimental Evaluations	40
3.2	Game-Theoretic Fairness Models	46
3.2.1	Incorporating Fairness into Agent Interactions Modeled as Two-Player Normal-Form Games	46
3.2.2	Incorporating Fairness into Infinitely Repeated Games with Conflicting Interests for Conflict Elimination	56
	References	69

4 Social Optimality in Cooperative Multiagent Systems	71
4.1 Reinforcement Social Learning of Coordination in Cooperative Games	72
4.1.1 Social Learning Framework	73
4.1.2 Experimental Evaluations	77
4.2 Reinforcement Social Learning of Coordination in General-Sum Games	82
4.2.1 Social Learning Framework	82
4.2.2 Analysis of the Learning Performance Under the Social Learning Framework	88
4.2.3 Experimental Evaluations	89
4.3 Achieving Socially Optimal Allocations Through Negotiation	100
4.3.1 Multiagent Resource Allocation Problem Through Negotiation	101
4.3.2 The APSOPA Protocol to Reach Socially Optimal Allocation	102
4.3.3 Convergence of APSOPA to Socially Optimal Allocation	108
4.3.4 Experimental Evaluation	110
References	112
5 Individual Rationality in Competitive Multiagent Systems	115
5.1 Introduction	115
5.2 Negotiation Model	117
5.3 <i>ABiNeS</i> : An Adaptive Bilateral Negotiating Strategy	119
5.3.1 Acceptance-Threshold (AT) Component	121
5.3.2 Next-Bid (NB) Component	122
5.3.3 Acceptance-Condition (AC) Component	124
5.3.4 Termination-Condition (TC) Component	125
5.4 Experimental Simulations and Evaluations	125
5.4.1 Experimental Settings	126
5.4.2 Experimental Results and Analysis: Efficiency	128
5.4.3 Detailed Analysis of <i>ABiNeS</i> Strategy	130
5.4.4 The Empirical Game-Theoretic Analysis: Robustness	133
5.5 Conclusion	141
References	141
6 Social Optimality in Competitive Multiagent Systems	143
6.1 Achieving Socially Optimal Solutions in the Context of Infinitely Repeated Games	143
6.1.1 Learning Environment and Goal	144
6.1.2 TaFSO: A Learning Approach Toward SOSNE Outcomes	147
6.1.3 Experimental Simulations	152
6.2 Achieving Socially Optimal Solutions in the Social Learning Framework	158
6.2.1 Social Learning Environment and Goal	159
6.2.2 Learning Framework	161

6.2.3 Experimental Simulations	164
References	169
7 Conclusion	171
Reference	173
A The 57 Structurally Distinct Games	175

Chapter 1

Introduction

Multiagent systems (MASs) have become a commonly adopted paradigm to model and solve real-world problems. Many competing definitions exist for a MAS. In this book, we consider a typical MAS as a system involving multiple autonomous software agents (or humans) interacting with each other with (possibly) conflicting interests and limited information, and the payoff of each agent is determined by the joint actions of all (or some) agents involved. Therefore, different from single-agent environments, in multiagent interaction environments, each agent needs to take other agents' behaviors into consideration when it makes its own decisions, since others' behaviors can directly influence what it expects from the system. The major question we seek to answer in this book can be summarized as follows: how can a *desirable goal* be achieved in different *multiagent interaction environments* where each agent may have its own limitations and (possibly) conflicting interests?

We divide different multiagent interaction environments into two major categories depending on the underlying intentional stance of the agents: cooperative multiagent environment and competitive multiagent environment. We say a multiagent interaction environment is cooperative if (most of) the agents within the system are willing to behave exactly as they are asked to even if this would cause conflicts with their individual interests. In other words, each agent is cooperative in achieving a commonly shared system-level goal even at the cost of its personal benefits. A multiagent interaction environment is considered to be competitive when each agent is only interested in maximizing its individual benefits. Thus, in competitive environments, agents may not have the incentive to follow what the system designer specifies and usually behave in purely individually rational manners.

We mainly focus on three major goals: fairness, social optimality, and individual rationality. To achieve either of the above goals, the research focus may vary and different approaches may be required depending on the nature of multiagent interaction problems being considered. In cooperative environments, since the agents are assumed to comply with the system designer's instructions, we usually need to consider how we can design an efficient learning strategy through which

the agents can effectively coordinate their behaviors toward the goal we target at (e.g., fairness, social optimality) subject to practical limitations (e.g., limited information available, limited computational, and communication ability). On the contrary, in competitive environments, we mainly focus on addressing the following two questions:

- How can an agent obtain as much individual benefits as possible in the presence of other selfish agents?
- The second one is from the perspective of mechanism design: how can selfish agents be incentivized to autonomously adjust their behaviors toward the goal (e.g., social optimality) we target at?

To model the strategic interactions in MASs, the most powerful and commonly used mathematical tool is game theory. It provides a rich set of powerful and convenient mathematical frameworks to model different forms of strategic interactions among agents. For example, normal-form games can be used to model the single-shot strategic interactions among agents in which each agent makes decision simultaneously; for finitely/infinately repeated games, which consists of finite/infinite repetition of stage games, they are naturally suitable for modeling the repeated (with unknown rounds) interactions among agents. Thus, in this book, we mainly study the problem of achieving different desirable goals by modeling agents' interactions as different types of games.

1.1 Overview of the Chapters

The book is organized in terms of the goals (fairness, social optimality, and individual rationality) we target at and the characteristic of the multiagent interaction environment (cooperative or competitive) we look at. Specifically, we focus on the following four parts, each of which corresponds to one chapter.

Fairness in cooperative MASs We pursue the goal of fairness in cooperative MASs from the following two perspectives. First, we describe an adaptive periodical strategy to enable the agents to coordinate toward fair and efficient outcomes in conflicting-interest games [1]. Second, two game-theoretic fairness models are described for two important multiagent interaction frameworks. The first is a game-theoretic fairness model for infinitely repeated game with conflict interest, in which it can be proved that there always exist fairness equilibria under which fair and optimal outcomes can be achieved [2]. The second is a game-theoretic fairness model for single-shot normal-form games, which can better explain actual human behaviors [3].

Toward social optimality in cooperative MASs In cooperative MASs, we describe two different types of learners for cooperative games, independent action learners and joint action learners, to coordinate on socially optimal outcomes under the social learning framework [4–6]. For general-sum games,

we introduce a learning framework under which the agents can effectively learn to coordinate on socially optimal outcomes by adaptively choosing between individual learning and social learning [7]. Last, we turn to look at the practical bilateral negotiation problem and introduce an efficient negotiation framework for agents to achieve socially optimal allocations in cooperative bilateral negotiation environments [8].

Maximizing individual benefits in competitive MASs In competitive MASs, we focus on the bilateral negotiation setting in which two agents negotiate with each other with the goal of maximizing their individual utilities through negotiation. We describe an adaptive negotiation *ABiNeS* strategy which empirically shows to be very effective in competing with the state-of-the-art negotiation strategies [9, 10]. It is worth mentioning that the implementation of this *ABiNeS* strategy wins the champion of the third international Automated Negotiating Agents Competition [11] in both the qualifying and final rounds.

Achieving social optimality in competitive MASs In competitive MASs, we, as the system designer, may be also interested in achieving social optimality. However, in competitive environments, the agents are individually rational and interested in maximizing their own utilities only, which may have conflicts with the goal of social optimality. To resolve this conflict, we introduce an incentive mechanism based on sequential play to enable rational agents to have the motivation to adopt the policy of achieving socially optimal outcomes, which are successfully applied in two different learning frameworks [12–14].

1.2 Guide to the Book

The intended reader of this book can be a graduate student or an advanced undergraduate or any researcher in the areas of computer science, game theory, and any other related areas. The material of each chapter is independently organized and the readers can have no difficulty of understanding the materials by jumping to any interested chapter freely. We outline the chapters that follow.

- Chapter 2 first reviews some background concepts in game theory which will be used throughout the book and then gives an overall review of previous work in multiagent learning literature, which are divided into two major categories: cooperative multiagent systems and competitive multiagent systems. The previous works within each category are further divided according to the solution concepts they adopt, and how they link with the works in this book is also mentioned.
- Chapter 3 focuses on investigating the question of how to achieve fairness within cooperative multiagent environments. We interpret fairness from the following two perspectives. First, we focus on the multiagent interaction scenarios with conflict interest and our goal is to achieve fairness among agents in terms of equalizing the payoffs of agents as much as possible. Second, our human beings are not purely selfish but care about fairness, which cannot be explained and

predicted by the classical game theory. Thus, it is worthwhile to develop fairness-based game-theoretic model to better explain and predict human behaviors. Overall, we investigate the goal of fairness in MAS from two perspectives: strategy design and fairness game-theoretic model design, which are introduced in Sects. 3.1 and 3.2, respectively.

- Chapter 4 turns to the investigation of how to achieve social optimality within cooperative multiagent environments. Achieving social optimality means maximizing the sum of all agents' payoff involved in the interaction. In cooperative multiagent environments, the agents are assumed to be willing to cooperate toward social optimality, but the limited information available to each agent might impede effective coordination toward socially optimal outcomes among them. We model the agent interactions using two different types of games, cooperative games and general-sum games, and describe the corresponding state-of-the-art social learning frameworks toward social optimality in Sects. 4.1 and 4.2, respectively. Lastly we consider a particular multiagent negotiation problem and describe an efficient negotiation mechanism for achieving socially optimal allocations in Sect. 4.3.
- Chapter 5 focuses on the problem of maximizing an agent's individual benefits within competitive multiagent systems. In competitive multiagent systems, each agent is usually only interested in maximizing its individual benefits, and a natural question from a single agent's perspective is how to design an efficient strategy to obtain as much payoff as possible by exploiting its partners during interactions. Therefore, in this chapter, we consider the competitive bilateral negotiation environment and introduce a state-of-the-art efficient negotiation strategy to maximize the individual utility of a negotiating agent against other competitive opponents.
- Chapter 6 is devoted to investigating how to achieve social optimality (maximizing system level's performance) within competitive multiagent systems. As we mentioned before, in competitive multiagent systems, the agents may not have the incentive to cooperate toward the goal of social optimality which may conflict with its individual benefits. Therefore, from the system designer's perspective, we need to design efficient mechanism to incentivize selfish agents to cooperate toward socially optimal outcomes. Therefore, in this chapter, we introduce an interesting variation of sequential play and show how this can be applied in different competitive multiagent interaction environments to facilitate agents to achieve socially optimal outcomes without undermining their individual rationality and autonomy, namely, two-agent repeated interaction framework (Sect. 6.1) and social learning framework (Sect. 6.2).
- Chapter 7 concludes the book with a summary of the book and also points out some future research directions for the readers.

References

1. Hao JY, Leung HF (2010) Strategy and fairness in repeated two-agent interaction. In: Proceedings of ICTAI'10, Arras, pp 3–6
2. Hao JY, Leung HF (2012) Incorporating fairness into infinitely repeated games with conflicting interests for conflicts elimination. In: Proceedings of ICTAI'12, Athens, pp 314–321
3. Hao JY, Leung HF (2012) Incorporating fairness into agent interactions modeled as two-player normal-form games. In: Proceedings of PRICAI'12, Kuching
4. Hao JY, Leung HF (2013) Reinforcement social learning of coordination in cooperative multi-agent systems(extended abstract). In: Proceedings of AAMAS'13, St. Paul, pp 1321–1322
5. Hao JY, Leung HF (2013) The dynamics of reinforcement social learning in cooperative multiagent systems. In: Proceedings of IJCAI'13, Beijing, pp 184–190
6. Hao JY, Leung HF, Ming Z (2014) Multiagent reinforcement social learning toward coordination in cooperative multiagent systems. *ACM Trans Auton Adapt Syst* 9(4):20
7. Hao JY, Leung HF (2011) Learning to achieve social rationality using tag mechanism in repeated interactions. In: Proceedings of ICTAI'11, Boca Raton, pp 148–155
8. Hao JY, Leung HF (2012) An efficient neogtiation protocol to achieve socially optimal allocation. In: PRIMA 2012: principles and practice of multi-agent systems, Kuching
9. Hao JY, Leung HF (2012) Abines: an adaptive bilateral negotiating strategy over multiple items. In: Proceedings of IAT'12, Macau, vol 2, pp 95–102
10. Hao JY, Song SZ, Leung HF, Ming Z (2014) An efficient and robust negotiating strategy in bilateral negotiations over multiple items. *Eng Appl Artif Intell* 34:45–57
11. Marsa-Maestre I, Lopez-Carmona MA, Ito T et al (2014) Novel insights in agent-based complex automated negotiation[M]. Springer, Tokyo
12. Hao JY, Leung HF (2012) Learning to achieve socially optimal solutions in general-sum games. In: PRICAI 2012: trends in artificial intelligence, Kuching, pp 88–99
13. Hao JY, Leung HF (2012) Achieving social optimality with influencer agents. In: Proceedings of complex'12, Santa Fe
14. Hao JY, Leung HF (2015) Introducing decision entrustment mechanism into repeated bilateral agent interactions to achieve social optimality. *Auton Agents Multi-Agent Syst* 29(4):658–682

Chapter 2

Background and Previous Work

In this chapter, we first review some concepts and terminologies from game theory and multiagent systems areas that will be used throughout the book in Sect. 2.1. After that, we explore previous work in multiagent learning literature by dividing them into two major categories: cooperative multiagent systems and competitive multiagent systems, distinguished by the underlying intentional stance of the agents within the system. Within each category, we review previous work in three different parts distinguished by the different goal that the work targets at. In Sect. 2.2, we focus on reviewing previous work whose goal falls into one of the following major solution concepts: Nash equilibrium, fairness, or social optimality. In Sect. 2.3, we focus on investigating previous work targeting at either of the following major solution concepts: Nash equilibrium, maximizing individual benefits, and Pareto-optimality.

2.1 Background

Game theory has been the most commonly adopted mathematical tool for us to model the strategic interactions among different agents in different interaction scenarios. We usually model the one-time interaction among agents as single-shot normal-form game and finite/infinite repeated games to model the repeated interaction among agents.

2.1.1 *Single-Shot Normal-Form Game*

In single-shot games, each agent is allowed to choose one action from its own action set simultaneously, and each agent will receive its own payoff based on the joint