

The London School of Economics and Political Science

Strategic Interdependence,
Hypothetical Bargaining, and Mutual
Advantage in Non-Cooperative Games

Mantas Radzvilas

A thesis submitted to the Department of Philosophy, Logic and Scientific Method
of the London School of Economics and Political Science
for the degree of Doctor of Philosophy, September 2016.

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any person is clearly identified in it). The copyright of this thesis rests with the author. Quotation from it is permitted, provided that full acknowledgement is made. This thesis may not be reproduced without my prior written consent. I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of 85,056 words.

Statement of conjoint work

I confirm that Chapter 2 is based on a paper co-authored with Jurgis Karpus (King's College London) and I contributed 50% of this work.

Mantas Radzvilas

Abstract

One of the conceptual limitations of the orthodox game theory is its inability to offer definitive theoretical predictions concerning the outcomes of non-cooperative games with multiple rationalizable outcomes. This prompted the emergence of goal-directed theories of reasoning – the team reasoning theory and the theory of hypothetical bargaining. Both theories suggest that people resolve non-cooperative games by using a reasoning algorithm which allows them to identify mutually advantageous solutions of non-cooperative games.

The primary aim of this thesis is to enrich the current debate on goal-directed reasoning theories by studying the extent to which the principles of the bargaining theory can be used to formally characterize the concept of mutual advantage in a way which is compatible with some of the conceptually compelling principles of orthodox game theory, such as individual rationality, incentive compatibility, and non-comparability of decision-makers' personal payoffs.

I discuss two formal characterizations of the concept of mutual advantage derived from the aforementioned goal-directed reasoning theories: A measure of mutual advantage developed in collaboration with Jurgis Karpus, which is broadly in line with the notion of mutual advantage suggested by Sugden (2011, 2015), and the benefit-equilibrating bargaining solution function, which is broadly in line with the principles underlying Conley and Wilkie's (2012) solution for Pareto optimal point selection problems with finite choice sets. I discuss the formal properties of each solution, as well as its theoretical predictions in a number of games. I also explore each solution concept's compatibility with orthodox game theory.

I also discuss the limitations of the aforementioned goal-directed reasoning theories. I argue that each theory offers a compelling explanation of how a certain *type* of decision-maker identifies the mutually advantageous solutions of non-cooperative games, but neither of them offers a definitive answer to the question of how people coordinate their actions in non-cooperative social interactions.

Contents

1	Introduction	11
1.1	The Basic Elements of Non-Cooperative Game Theory . . .	14
1.1.1	A Formal Representation of a Normal Form Game . .	14
1.1.2	What do Payoffs Actually Represent?	15
1.1.3	The Complete Information Assumption	19
1.1.4	Best Response	22
1.1.5	Best Response and The Nash equilibrium	25
1.1.6	Rationality as Best-Response Reasoning	29
1.2	Social Coordination and the Conceptual Limitations of the Best-Response Reasoning Models	32
1.3	Social Coordination Theories	37
1.3.1	Cognitive Hierarchy Theory	37
1.3.2	Stackelberg Reasoning	42
1.3.3	Social Conventions Theory	43
1.4	Strategic Reasoning and Mutual Advantage	48
1.4.1	Coalitional Rationalizability	48
1.4.2	Goal-Directed Reasoning and Mutual Advantage . . .	52
1.4.3	Team Reasoning and Mutual Advantage	53
1.4.4	Hypothetical Bargaining	55
1.5	The Structure of the Thesis	59
2	Team Reasoning and the Measure of Mutual Advantage	60
2.1	Introduction	60
2.2	Team Reasoning	63
2.2.1	What is Team Reasoning?	63
2.2.2	Team Reasoning and Transformations of Payoffs . . .	65
2.3	Team Interests and the Notion of Mutual Advantage	68
2.3.1	Self-Sacrifice and Mutual Advantage	68
2.3.2	Interpersonal Comparisons of Payoffs	69
2.4	Team Interests as the Maximization of Mutual Advantage .	71
2.4.1	Formalization	73

2.4.2	Two Properties of the Function	74
2.4.3	Reference points	75
2.4.4	Examples	77
2.4.5	No Irrelevant Player	79
2.4.6	Independence of Irrelevant Strategies	80
2.4.7	Interpersonal Comparisons of Advantage	81
2.5	Mutual Advantage and the Problem of Coordination	85
2.6	The Triggers of Team Reasoning	88
2.7	Conceptual Limitations	91
2.7.1	Implicit Motivation Transformations	91
2.7.2	Stability Issues	94
2.8	Conclusion	95
3	Hypothetical Bargaining and the Equilibrium Selection In Non-Cooperative Games	98
3.1	Introduction	98
3.2	Hypothetical Bargaining	104
3.2.1	Misyak and Chater's Virtual Bargaining Model	104
3.2.2	The Limitations of the Model	106
3.3	The Ordinal Benefit-Equilibrating Solution	111
3.3.1	The Intuition Behind the Ordinal BE Solution	111
3.3.2	Formalization	116
3.3.3	Ordinal BE Solution Properties	119
3.3.4	Examples	122
3.4	The Cardinal Benefit-Equilibrating Solution	126
3.4.1	The Intuition Behind the Cardinal BE Solution	126
3.4.2	Formalization	131
3.4.3	The Properties of the Cardinal BE Solution	133
3.4.4	Relation to other Bargaining Solutions	135
3.4.5	Examples	137
3.4.6	Hypothetical Bargaining in Pie Games	140
3.5	Possible Extensions of the Model	144
3.5.1	Application to N -Player Games	144
3.5.2	Social Interactions and the Set of Feasible Agreements	150
3.6	Conclusion	152
4	Hypothetical Bargaining, Social Coordination and Evolu- tion	155
4.1	Introduction	155
4.2	Hypothetical Bargaining and Rational Choice	160
4.2.1	Hypothetical Bargaining as a Reasoning Algorithm	160

4.2.2	The Rationalization Problem	163
4.3	Hypothetical Bargaining and Social Coordination	175
4.3.1	Hypothetical Bargaining and the Problem of Common Beliefs	175
4.4	Hypothetical Bargaining and Evolution	180
4.4.1	Evolutionary Game Theory and Population Dynamics	180
4.4.2	Hypothetical Bargaining as an Evolutionary Adaptation	187
4.5	The Non-Uniqueness Problem	200
4.6	Conclusion	202
5	Conclusion	205
	Bibliography	212

List of Figures

1.1	Prisoner's Dilemma	17
1.2	Weak Dominance game	31
1.3	Coordination game	33
1.4	Hi-Lo game	35
1.5	Coordination game	40
1.6	Battle of the Sexes game	45
1.7	'Weak' Common Interest game	50
2.1	Hi-Lo game	60
2.2	Prisoner's Dilemma game	64
2.3	Monetary Prisoner's Dilemma game (a) and its transformation (b)	65
2.4	A variation of the Prisoner's Dilemma game	69
2.5	Another representation of preferences in the Prisoner's Dilemma game.	70
2.6	Chicken (a) and High Maximin (b) games	77
2.7	Three versions of the extended Hi-Lo game	85
2.8	Coordination game with a conflict of players' preferences	93
3.1	Common interest game	99
3.2	Hi-Lo game	104
3.3	4x4 games with three weakly Pareto optimal Nash equilibria	108
3.4	Ordinal coordination game	112
3.5	Ordinal coordination game with a unique Pareto efficient Nash equilibrium	122
3.6	Ordinal Chicken game	123
3.7	Divide-the-Cake game (even number of pieces)	124
3.8	Divide-the-Cake game (odd number of pieces)	125
3.9	Coordination game with three weakly Pareto optimal outcomes	126
3.10	Coordination game with three weakly Pareto optimal outcomes	138
3.11	Chicken game	139
3.12	Battle of the Sexes game	139

3.13	An example of a pie game from Faillo et al. (2013, 2016) experiments (Source: Faillo et al. 2013)	141
3.14	An example of a pie game depicted in Figure 3.13 as seen by the two interacting players in one treatment. The positions of the three slices were varied across treatments (Source: Faillo et al. 2013)	141
3.15	3x3 pie game depicted in Figure 22 represented in normal form	142
3.16	Three player coordination game played by three hypothetical bargainers	146
3.17	Three player coordination game, in which the matrix player is not a hypothetical bargainer	149
3.18	The Boobytrap Game	151
3.19	Two versions of the extended Hi-Lo game	153
4.1	Common interest game	156
4.2	Coordination game	176
4.3	Coordination game	192
4.4	Mixed motive game	195
4.5	Mixed motive game under (a) replicator dynamics and (b) best response dynamics	196
4.6	Mixed motive game with efficient exploiters	197
4.7	Mixed motive game with efficient exploiters under (a) best response dynamics and (b) replicator dynamics.	198
4.8	Extended Hi-Lo game	201

List of Tables

2.1	The levels of individual and mutual advantage associated with each outcome.	78
2.2	A summary of the results for the remaining games	78
3.1	Players' foregone preferred alternatives	113
3.2	Comparison of players' losses of attainable individual advantage	123
3.3	Comparison of players' losses of attainable individual advantage	123
3.4	Comparison of players' losses of attainable individual advantage	125
3.5	Comparison of players' losses of attainable individual advantage	125
3.6	Comparison of players' losses of maximum individual advantage	138
3.7	Comparison of players' losses of maximum individual advantage	139
3.8	Comparison of players' losses of maximum individual advantage	140
3.9	Summary of Faillo et al 2013 results. Choices predicted by the cardinal BE model are indicated by be . Choices predicted by the team reasoning model are indicated by τ	142
3.10	Summary of Crawford et al 2010 results. The choices of player 1 (P1) and player 2 (P2) are presented separately. Choices predicted by the cardinal BE model are indicated by be . Choices predicted by the team reasoning model are indicated by τ	143
3.11	Players' levels of individual advantage, losses of maximum individual advantage, and differences of individual advantage losses	146
3.12	The ratios of each players individual advantage level to the sum of players individual advantage levels, and the distance between each players actual ratio and the strictly egalitarian ratio	148

Acknowledgments

Several people have helped me to improve this thesis. I should first mention my primary supervisor, Jason McKenzie Alexander. Without his open mindedness to my ideas about game theory, profound knowledge of game theory and mathematics, as well as his exceptional attention to detail, my thesis would have been considerably poorer. I am also indebted to my secondary supervisor, Richard Bradley, whose critical yet encouraging comments and unwavering support gave me the intellectual courage and determination to develop my ideas about game theory. I would like to thank them both for all the support they provided me with throughout these years.

I would like to thank my examiners, Robert Sugden and Francesco Guala, with whom I had an intellectually challenging, yet ultimately very enjoyable and educational exchange about the topics of this dissertation. I would also like to thank Alex Voorhoeve, whose support and critical feedback on my work was invaluable.

I am also grateful to my very good friends and colleagues, Jurgis Karpus and Nicolas Wüthrich, whose feedback on numerous drafts of my work was extremely helpful.

My fellow PhD students James Nguyen, Goretí Faria, Deren Olgun, Johannes Himmelreich, Catherine M. Greene, Aron Vallinder, Philippe van Basshuysen and Tom Rowe have been great friends and companions in moments of joy and distress during my life as a PhD student. Special thanks go to my good friend Susanne Burri, whose advices and encouragement were of great help to me during my first year as a PhD student.

I thank my partner, Alina Mickevic, for her unwavering patience, understanding and loving support, especially during the last year of my PhD.

Most importantly, I am profoundly grateful to my parents, Aldona and Vytautas, who encouraged me to pursue what I am truly passionate about. Their unconditional love and support sustained me throughout my PhD and proved vital for the completion of this thesis.

Chapter 1

Introduction

In most general terms, orthodox game theory¹ can be defined as ‘the study of mathematical models of conflict and cooperation between intelligent rational decision-makers’ (Myerson 1991: 1). In terms of research objectives, game theory can be divided into normative and descriptive branches: Normative game theory explores the ‘nature and the consequences of idealized full rationality in strategic interactions’, while descriptive game theory ‘aims at the explanation and prediction of observed behavior’. (Selten 1988: vii). In other words, normative game theory focuses on answering the question of what a perfectly rational decision-maker *should* be expected to do in certain idealized strategic situations, while the aim of descriptive game theory is to develop game theoretic models which would better explain and/or predict people’s behaviour in certain types of real-world interdependent decision problems. According to the proponents of the position known as ‘methodological dualism’, game theorists both can and should avoid a conflation of normative and descriptive research goals (see, for example, Aumann 1985).

As has been pointed out by Selten (1988), ‘the distinction between normative and descriptive game theory is blurred in the practice of applied research’, since the ‘methods developed in normative theory are used in the

¹In orthodox game theory, each agent is assumed to be a rational decision-maker who engages in strategic deliberations aimed at finding a course of action which maximally advances his or her personal interests. In some evolutionary game theory models, however, players need not be thinking creatures at all. For example, the models of replicator dynamics are based on assumption that each player is simply programmed to play a particular strategy in every interaction with other individuals, irrespective of what strategies the other players play. This chapter is dedicated primarily to the exploration of the conceptual problems of ‘rationalistic’ orthodox game theory. Some of the issues pertaining to conceptual foundations of evolutionary game theory will be discussed in chapter 4. For a more detailed discussion of the differences between rationalistic and evolutionary branches of game theory, see, for example, Binmore 2008.

analysis of applied models in the hope for empirical relevance.’ (Selten 1988: vii). In other words, analytic methods and solution concepts developed for the purposes of normative analysis of highly idealized mathematical models of strategic interactions are also used for the development and analysis of models which purport to explain people’s actions in certain types of social interactions.

Due to its abstract nature and, consequently, extremely general scope, game theory has become one of the most important tools of social scientists: Various more or less complex interactions between two or more social agents are modelled as games played by rational agents. The theoretical predictions of decision-makers’ actions in idealized mathematical models of strategic interactions, or *games*, are viewed as offering ‘insights into any economic, political, or social situation that involves individuals who have different goals and preferences.’ (Myerson 1991: xi). Yet despite its widespread use, orthodox game theoretic analysis has certain conceptual limitations: Even the simplest of games have multiple rational solutions. From the perspective of orthodox game theory, every rational solution of a game is as valid and credible as any other. This leads to a problem which has been clearly stated by Bacharach and Bernasconi:

‘It has become apparent that for many important classes of games traditional game theory is indeterminate, since tightening traditional solution concepts to the limit still leaves multiple solutions. Explaining players’ behaviour therefore requires an addition to game theory, a theory of how players select one solution from several that are equally eligible as far as game theory is concerned.’ (Bacharach and Bernasconi 1997: 1-2)

In other words, one of the widely recognized conceptual limitations of the orthodox game theory is its inability to offer theoretical predictions of players’ actions in games with multiple rational solutions. This is a non-trivial problem, since many real-world interdependent decision problems are modelled as non-cooperative games with multiple rational solutions. Therefore, descriptive models based on the principles of orthodox game theory cannot explain and/or predict social agents’ behaviour in many important real-world social interactions, thus making descriptive game theory a far less useful explanatory tool than it could *prima facie* be expected to be.

This prompted the emergence of a number of more or less empirically successful descriptive theories, all of which purport to explain how people choose their actions in games with multiple rational solutions. Two of the more recent theories are the team reasoning theory, pioneered by Sugden (1993) and the hypothetical bargaining theory, pioneered by Misyak and Chater (2014)

and Misyak et al. (2014). Both theories suggest that people resolve non-cooperative games by identifying outcomes which are mutually advantageous – individually advantageous for every interacting decision-maker. By identifying and playing their part in the attainment of mutually advantageous outcomes, decision-makers either manage to resolve the coordination problem completely or, in games with multiple mutually advantageous outcomes, increase the coordination success rate. Both the team reasoning theory and the hypothetical bargaining theory can be viewed as goal-directed reasoning theories: In both theories, decision-makers are assumed to follow a specific reasoning procedure in order to identify a combination of strategies, or *strategy profile*, that leads to the attainment of a specific goal – a mutually advantageous outcome of the game.

A number of properties have been discussed in the literature for a mutually advantageous outcome to satisfy: Pareto efficiency, feasibility, successful coordination of interacting decision-makers' actions and equitable distribution of individuals' personal payoff gains. Yet so far very few formal characterizations of the concept of mutual advantage which could be incorporated into formal game theoretic analysis have been proposed.

The primary aim of this thesis is to enrich the current debate on goal-directed reasoning models with a study of how the formal concept of mutual advantage could be incorporated into formal game theoretic analysis, as well as to explore the extent to which such models of strategic reasoning are compatible with the principles of orthodox game theory. In this study, two novel formal characterizations of the concept of mutual advantage, which can be derived from the principles of team reasoning theory and hypothetical bargaining theory, will be proposed: A measure of mutual advantage developed in collaboration with Karpus (Karpus and Radzvilas 2016), which is broadly in line with the principles of the version of team reasoning theory suggested by Sugden (2011, 2015), and a benefit-equilibrating (BE) solution concept, which is broadly in line with the principles of hypothetical bargaining theory suggested by Misyak and Chater (2014) and Misyak et al. (2014). The formal properties as well as the theoretical predictions of each solution concept will be discussed. In addition, each solution concept's compatibility with the principles of orthodox game theory will be explored.

In section 1 of this introductory chapter, I introduce the basic concepts of the orthodox non-cooperative game theory, such as game, payoff, strategy, best response, best-response reasoning and the Nash equilibrium. In section 2 I will discuss the best-response reasoning model and explain the reasons of why it cannot rule out certain intuitively unreasonable solutions of non-cooperative games. In section 3 I will briefly overview some of the descriptive

theories which purport to explain how people resolve games with multiple rational solutions. In section 4 I will introduce descriptive theories which purport to explain people's behaviour as resulting from their attempts to resolve games in a mutually advantageous way. With section 5 I conclude this introductory chapter with an outline of the structure of this thesis.

1.1 The Basic Elements of Non-Cooperative Game Theory

A game is a formal representation of a certain type of interdependent decision problem – a complete formal description of the strategic interaction, which includes all of the constraints on the actions that individuals *can* take as well as players' personal interests, but does not include a specification of actions that individuals *do* take (Rubinstein and Osborne 1994). The formal structure of the game represents both the structure of the real-world strategic interaction – the sets of strategies available to each of the interacting players in a particular interdependent decision problem – and players' preferences over the physical outcomes of the game – the combinations of players' actions, resulting from each player's strategy choice (for extensive discussion, see, for example, Luce and Raiffa 1957, Rubinstein 1991 and Binmore 2009a).

1.1.1 A Formal Representation of a Normal Form Game

Formally, a normal form game² Γ can be defined as a triple $(I, \{S_i, u_i\}_{i \in I})$, where $I = \{1, \dots, m\}$ is the set of players of the game, $S_i = \{1, 2, \dots, k_i\}$ is the set of pure *strategies* of every player $i \in I$, and $u_i : \mathbf{S} \rightarrow \mathbb{R}$ is the payoff function of each player $i \in I$, where $\mathbf{S} = \times_{i \in I} S_i$ is the set of *strategy profiles*, or *outcomes*, of Γ .

A *strategy* is a complete algorithm for playing the game, which fully specifies what the player does (i.e. what action or actions the player takes) in every possible situation throughout the game (for extensive discussion, see Rubinstein 1991)³. A *strategy profile* $\mathbf{s} = (s_1, \dots, s_m)$, where $s_i \in S_i$ is a pure strategy of player $i \in I$, is a vector of pure strategies, which fully specifies players' actions in the game. The set of strategy profiles of the game defines the set of all the possible outcomes.

²The following chapters will focus on the analysis of players' behaviour in one-shot complete information games, which can be represented as normal form games.

³In terms of epistemic game theory, player's strategy is a function from states of the world to actions, where each state of the world is characterized by a specific combination of all the other players' strategy choices. For extensive discussion, see Perea 2012.

A payoff function $u_i : \mathbf{S} \rightarrow \mathbb{R}$ of each player $i \in I$ maps all the possible strategy profiles of the game Γ into the set of real numbers. In standard game theoretic analysis of non-cooperative normal form games, it is assumed that players have complete preferences over the possible outcomes, and so each player $i \in I$ has a certain payoff $u_i(\mathbf{s}) \in \mathbb{R}$ associated with every possible pure strategy profile $\mathbf{s} \in \mathbf{S}$. The combined pure strategy payoff function $u : \mathbf{S} \rightarrow \mathbb{R}$ of Γ assigns a full vector $u(\mathbf{s}) = (u_1(\mathbf{s}), \dots, u_m(\mathbf{s}))$ of payoffs to every strategy profile $\mathbf{s} \in \mathbf{S}$.

1.1.2 What do Payoffs Actually Represent?

An important question pertaining to players' payoff functions is what exactly the real numbers associated with game's outcomes actually represent. The orthodox interpretation of payoffs is that they represent players' von Neumann and Morgenstern utilities. In rational decision theory, von Neumann and Morgenstern utilities represent players' preferences over the choice options, satisfying the axioms of the expected utility theory – *completeness*, *transitivity*, *independence* and *continuity*⁴. These axioms are viewed as defining choice consistency, and must be satisfied in order for the construction of the von Neumann and Morgenstern utility function to be possible (for a detailed discussion of the expected utility theory, see, for example, Luce and Raiffa 1957 and Kreps 1988).

Decision-maker's von Neumann and Morgenstern utility function should capture all the motivations relevant for a decision-maker's choice in a particular decision problem, including his or her attitude to risk. According to the revealed preference theory, pioneered by Samuelson (1938) and later advocated by Little (1949) and many other economists, decision-makers' pref-

⁴Let x , y and z be three alternative choice options. Let $>$, \sim and \geq denote strong preference, indifference and weak preference relations over choice options respectively. The choice consistency axioms are the following:

1. *Completeness*: For any two choice options x and y , it is always the case that $x \geq y$ or $y \geq x$.
2. *Transitivity*: For any three choice options x , y and z , if $x \geq y$ and $y \geq z$, then $x \geq z$.
3. *Independence*: For any three choice options x , y and z , if $x \geq y$, then, for any probability $p \in [0, 1]$, it must be the case that $px + (1 - p)z \geq py + (1 - p)z$.
4. *Continuity*: For any three choice options x , y and z , if $x \geq y \geq z$, then there must exist a probability $p \in [0, 1]$ such that $px + (1 - p)z \sim y$.

It can be shown that if decision maker's preferences satisfy axioms 1-4, they can be represented numerically by a function u , such that $x \geq y$ if and only if $u(x) \geq u(y)$. For technical details and proofs, see, for example, Kreps 1988.

erences over choice options are revealed in their choices, and so a decision-maker's preferential ranking of the choice options can be reconstructed from the observations of decision maker's choice behaviour (see, for example, Samuelson 1938 and Little 1949). The von Neumann and Morgenstern utility function can, at least in theory, be constructed from the observations of decision-maker's choices between objective lotteries (that is, lotteries involving some randomization device, the workings of which are known to the decision-maker) over choice options. Decision maker's choices between lotteries should, in theory, reveal his or her attitude to risk (for extensive discussion, see Luce and Raiffa 1957, Kreps 1988 and Hausman 2012). The resulting numerical representation of decision maker's preferences should be the von Neumann and Morgenstern utility function, which is unique only up to positive affine transformations: If u is a function representing decision maker's preferences over choice options, then so is any function $u' = au + c$, where $a > 0$ and c are constants (for a detailed discussion of why this is so see, for example, Luce and Raiffa 1957).

The interpersonal non-comparability of utility is a standard assumption of rational choice theory. From the axioms of the expected utility theory, one cannot draw a conclusion that the interpersonal comparisons of numbers representing decision-makers' cardinal preferences are meaningful. It must be emphasized, however, that the expected utility theory does not negate the possibility that interpersonal comparisons of utilities could be shown to be meaningful. That is, it does not show that a theory providing conceptual tools which would allow for meaningful interpersonal comparisons of utility numbers could not be developed⁵. Yet the expected utility theory itself does not justify the interpersonal comparisons of utility. In the absence of a compelling theory of how meaningful interpersonal comparisons of utility could be made, most of the orthodox decision and game theorists take the utility numbers not to be interpersonally comparable.

In game theory, players' payoffs do not represent their preferences over the available choice options – their strategies. The payoffs represent players' preferences over the *outcomes* of the game. An important principle of game theory, advocated by Binmore (1992, 2005, 2009a,b) and many other game theorists, is that the payoff structure of the game has to fully capture everything that is motivationally relevant in players' evaluations of the possible outcomes of the game. This principle implies that payoffs must capture all the relevant motivations of players, no matter what the nature of

⁵For a defense of a theoretical position that such a theory of interpersonal comparisons of utilities is conceptually possible, see, for example, Binmore 2005. For a critical overview of the attempts to justify interpersonal comparisons of utility, see Hammond 1991.

those motivations is, and that game theoretic analysis of any real-world interdependent decision problem is only meaningful if the payoffs of the game accurately represent the motivations of the interacting social agents (for a comprehensive discussion of this principle, see, for example, Binmore 1992, 2005 and 2009a).

An important problem associated with this definition of payoffs is that players' of the game choose strategies, not outcomes. An outcome, which is determined by a combination of players' strategy choices, is not something that any of the players can individually choose. As has been pointed out by Hausman (2012), players' strategy choices reflect not only their preferences over the outcomes of the game, but also their *beliefs* about the opponents' strategy choices. This means that players' strategy choices do not reveal, in Hausman's terms, the 'all-things-considered' preferential rankings of outcomes. They reveal players' preferential rankings of the available strategies, based both on players' preferences over outcomes and their beliefs about the opponents' strategy choices – beliefs, which should play no role in players' evaluations of outcomes (see Hausman 2012).

Binmore's (2009a) response to this problem is that players' preferences over the outcomes could, in principle, be elicited with 'games against nature' – one person decision problems, in which one of the players has to choose a strategy while already knowing the combination of opponents' strategy choices. From the observations of player's strategy choices in decision problems representing all the possible combinations of opponents' strategy choices, it should be possible to reconstruct player's preferential ranking of outcomes (see Binmore 2009a).

The problem with this method is that a rational player may never choose to bring about certain outcomes in one player games, which means that player's complete preferential ranking of all the possible outcomes of the game will not be revealed. For example, consider the Prisoner's Dilemma game. The game is shown in Figure 1.1, where one player chooses between the two options identified by rows and the other—by columns. The left and the right number in each cell represents row and column player's payoffs respectively.

	<i>c</i>	<i>d</i>
<i>c</i>	3, 3	0, 4
<i>d</i>	4, 0	1, 1

Figure 1.1: Prisoner's Dilemma

This game has a unique pure strategy Nash equilibrium (d, d) . Strategy c is, for every player, strictly dominated by strategy d . This means that it is never optimal to choose c , no matter what the opponent does. According to orthodox game theory, a rational player should never choose strategy c .

Suppose that two rational social agents are playing a Prisoner's Dilemma game. The observer does not know the structure of individuals' motivations, and so attempts to reconstruct each individual's preferential ranking of outcomes by observing his or her choices in two 'games against nature'. According to Binmore's suggestion, in one game an individual should be given a choice between strategies c and d while knowing that the other individual has chosen strategy c . In the second game, an individual should be given a choice between strategies c and d while knowing that the other individual has chosen strategy d . Assuming that the choosing individual is rational and that his or her personal motivations are like the motivations of a player in the Prisoner's Dilemma game (Figure 1.1), s/he should be observed choosing strategy d in both 'games against nature'. The observer could conclude that individual prefers outcome (d, c) over outcome (c, c) , and outcome (d, d) over outcome (c, d) . Individual's choices, however, reveal neither the preference relation between outcomes (c, d) and (d, c) , nor the preference relation between outcomes (d, d) and (c, c) . Therefore, a complete preferential ranking of all the outcomes of the game cannot be reconstructed, and the payoff matrix like the one depicted in Figure 1.1 cannot be produced (for an extensive discussion of the limitations of the revealed preference approach, see Rubinstein and Salant 2008 and Hausman 2012).

The preceding example suggests that individual's preferences over the outcomes could only be reconstructed from his or her choices between options in a set of hypothetical decision problems, each of which gives the decision-maker a choice between a pair of outcomes of the game. Individual's choices in a set of decision problems representing every possible pair of outcomes of a game should, in theory, give the observer enough information to construct decision-maker's complete *ordinal* preferential ranking of outcomes. This information, however, would not be sufficient for the von Neumann and Morgenstern utility representation of decision-maker's preferences. To capture decision-maker's attitude to risk, his or her choices between objective lotteries over the outcomes of the game would have to be observed. Since the outcomes of many real-world interdependent decision problems cannot be expressed in terms of quantities of material resources, the elicitation of real-world players' preferences is a challenging problem. Currently no compelling answer to this problem can be found in game theoretic literature, and so an accurate representation of people's cardinal preferences over the

outcomes of real-world interdependent decision problem can be viewed as one of the major methodological challenges of descriptive game theory.

In standard game theoretic models of complete information, the payoff structure of the game is assumed to be common knowledge. That is, it is assumed that each player knows the payoff structure of the game, knows that every other player knows it, knows that every player knows that every player knows it, and so on *ad infinitum*. In many real-world interdependent decision problems, the common knowledge of cardinal payoffs assumption is simply unrealistic: In many cases people, at best, know each other's ordinal preferences over outcomes. Therefore, a solution concept which can be defined in terms of purely ordinal information about players' preferences may in some cases be better suited to explain how they resolve an interdependent decision problem. For this reason, a solution concept which can be defined on the basis of ordinal information about players' preferences will be given considerable attention in the following chapters.

1.1.3 The Complete Information Assumption

In standard game theoretic analysis of non-cooperative complete information games, the payoff structure of the game is assumed to be common knowledge. In other words, it is assumed that each player's cardinal preferences over outcomes and the set of available strategies are common knowledge among the interacting players. It plays an important role in game theoretic analysis of one-shot games. If the payoff structure of the game were not common knowledge, the players would face an incomplete information game where uncertainty would be extreme: Each player would have to consider a potentially infinite set of games that the opponents could be playing, and assign a subjective probability distribution over it. In addition, the player would have to assign subjective probabilities to opponents' strategies in every possible game, and then choose an optimal response (for an extensive discussion of incomplete information games see, for example, Kreps 1990 and Fudenberg and Tirole 1991). The structure of the resulting complex incomplete information game would not be the same as the structure of the original one-shot game, and so the conclusions resulting from game theoretic analysis of players' strategy choices in the incomplete information game could not be taken to be representative of the conclusions of game theoretic analysis of players' strategy choices in the commonly known one-shot game.

In normative game theory, the complete information assumption is unproblematic: Normative game theoretic analysis is simply a theoretical exploration of how perfectly rational agents should behave in various highly idealized strategic interactions, some of which may not even resemble any

real-world social interactions. Its status in descriptive applications of game theoretic models is, however, far more ambiguous. First, complete information assumption implies that every social agent somehow knows all the relevant motivations of every interacting individual, even the subtle psychological and social motivations that influence individuals' evaluations of outcomes, such as his or her attitude to risk and inequity, sensitivity to social norms and personal moral considerations. Even very close family members cannot be expected to be always fully aware of each other's motivations, and such awareness seems to be even less likely to be present in social interactions involving less closely related social agents.

Second, the complete information assumption implies that each interacting individual is certain that every other interacting individual understands the structure of the interdependent decision problem in the same way as s/he does. Social reality, however, is complex: Social agents continuously engage in new or repeated interactions with other agents, and most social agents engage in more than one type of social interaction at once. The boundaries between different types of interactions are often less than obvious. Social agent's ability to recognize social interaction's type often depends on his or her ability to interpret various subtle situational cues and signals, which in turn depends on social agent's knowledge of other players' motivations, as well as knowledge of an intricate network of cultural norms and conventions (for a detailed discussion, see Bicchieri 2006 and Gintis 2008). Even if players were perfectly rational Bayesian deliberators, their ability to identify different types of social interactions would often rely on their background knowledge of social rules and practices. Therefore, common knowledge of rationality alone does not give social agents a reason to expect each other to understand the structure of the social interaction in the same way. The players who share no substantial information about each other's background and have not engaged in repeated interactions with each other are likely to be uncertain about each other's ability to correctly identify the type of social interaction. The complete information assumption applied in game theoretic explanations of social behaviour may, in a considerable number of cases, be a false assumption about social agents' knowledge of the game and its players.

The extent to which the use of such a seemingly unrealistic assumption in a descriptive model is problematic largely depends on the interpretation of the explanatory scope of the model. Infante et al. (2014) suggest two possible interpretations of behavioural models. Under the first interpretation, descriptive models are approximately accurate descriptions of observable choices. That is, a descriptive model only suggests that people behave *as if* they were expected utility maximizers with rational preferences over

outcomes (i.e. consistent and stable preferences which represent decision-maker's all-things-considered evaluations of possible outcomes). The aim of the model is not to explain how people actually reason, just to predict their actions. It does not offer an answer to the question of whether people's preferences satisfy the axioms of the expected utility theory. Under this interpretation, the complete information assumption plays merely an instrumental role: Social agents are assumed to behave *as if* the information about the structure of the game were common knowledge. The model does not suggest that social agents *actually* know each others payoffs and available strategies, know that every other player knows it, and so on *ad infinitum*. Under this 'thin' interpretation of descriptive models, the complete information assumption seems to be relatively unproblematic: The model explains the behaviour of social agents by describing and/or predicting it, not by providing an explanation of why social agents choose one or another action.

Another possible interpretation is that descriptive models offer an approximately true description of how people reason. That is, a descriptive model is an approximately accurate description of the process of reasoning by which people arrive at their strategy choices. Under this interpretation of descriptive models, the complete information assumption plays a non-trivial role in the explanation of people's choices: If the process of reasoning by which people arrive at their choices is (roughly) similar to the reasoning algorithm which underlies the orthodox game theoretic analysis of strategy choices, then player's beliefs obviously play an important role in the explanation of why s/he has identified a particular action as an optimal response to a particular interdependent decision problem. A question of whether the structure of a social interaction can reasonably be assumed to be common knowledge among the interacting players must play an important role in the evaluation of the explanatory relevance of the suggested model.

If behavioural models were truly intended to be used merely as tools for description and/or prediction of observed behaviour, the complete information assumption could be deemed completely unproblematic. The actual practices of behavioural economists, however, show this not to be the case: A clear distinction between the two interpretations of behavioural models is not maintained in the practice of applied research. Game theorists working on behavioural models often put considerable effort in providing empirical justifications for the assumptions pertaining to social agents' preferences and beliefs. Such behaviour should be taken as an indication that descriptive models are, at least implicitly, treated as approximately true descriptions

of the process of reasoning by which people arrive at their action choices⁶. Because of this, the information assumptions of descriptive models will not be treated as merely instrumentally useful components of the models, but rather as approximately true descriptions of social agents' beliefs. The problems associated with the epistemic assumptions of descriptive theories will play an important role in the following discussion.

1.1.4 Best Response

The original definition of best response, due to Nash (1950b, 1951), is that best response is a pure or mixed strategy which maximizes player's expected payoff against a fixed combinations of opponents' strategies.

This idea can be easily represented as a formal concept. Suppose that players are playing a normal form game $\Gamma = (I, \{S_i, u_i\}_{i \in I})$, where $I = \{1, \dots, m\}$ is the set of players, S_i is a set of pure strategies of player $i \in I$, and $u_i : \mathbf{S} \rightarrow \mathbb{R}$ is the von Neumann and Morgenstern utility function, where $\mathbf{S} = \times_{i \in I} S_i$ is the set of strategy profiles, or outcomes, of Γ . A mixed strategy of player $i \in I$ is a probability distribution over S_i . Let Σ_i denote a set of all such probability distributions and let $\sigma_i \in \Sigma_i$ denote a mixed strategy of $i \in I$, where $\sigma_i(s_i)$ is a probability assigned to pure strategy $s_i \in S_i$. A mixed strategy outcome is a mixed strategy profile $\sigma = (\sigma_1, \dots, \sigma_m)$. Let $\Sigma = \times_{i \in I} \Sigma_i$ denote the set of all mixed strategy profiles of Γ . The expected payoff associated with a mixed strategy profile $\sigma \in \Sigma$ is, for every player $i \in I$,

$$u_i(\sigma) = \sum_{\mathbf{s} \in \mathbf{S}} \left(\prod_{i \in I} \sigma_i(s_i) \right) u_i(\mathbf{s}). \quad (1.1)$$

Let $\sigma_{-i} = (\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_m)$ be a combination of mixed strategies of all the players other than $i \in I$. Strategy $\sigma_i \in \Sigma_i$ is player i 's best response to

⁶As has been noted by Infante et al. (2014), in recent years there has been a trend among behavioural economists to appeal to dual-process theories of the mind in order to justify the assumption that real-world decision-makers have rational preferences – stable and cognitive bias-free preferences which are consistent with the axioms of the expected utility theory. This justification strategy rests on assumption that system 2 is a general-purpose mode of practical reasoning which produces preferences and modes of strategic reasoning consistent with the expected utility maximization model (for details, see Infante et al. 2016). A similar trend can be noticed among behavioural models suggested by game theorists: System 2 is effectively assumed to be a best-response reasoning mode. For example, a belief-based model of homophily, suggested by Kets and Sandroni, rests on assumption that system 2 is a reasoning mode which, if it gets activated, ‘overrules’ the choice impulses generated by system 1, unless following the impulse in a particular game is a best response (for details, see Kets and Sandroni 2014).

σ_{-i} if it is the case that

$$u_i(\sigma_i, \sigma_{-i}) \geq u_i(\tilde{\sigma}_i, \sigma_{-i}) \quad \forall \tilde{\sigma}_i \in \Sigma_i. \quad (1.2)$$

Let $\beta_i(\sigma_{-i}) \subseteq \Sigma_i$ be a set of i 's best responses to a combination of opponents' strategies σ_{-i} . It can be defined as follows:

$$\beta_i(\sigma_{-i}) = \{\sigma_i \in \Sigma_i : u_i(\sigma_i, \sigma_{-i}) \geq u_i(\tilde{\sigma}_i, \sigma_{-i}) \quad \forall \tilde{\sigma}_i \in \Sigma_i\}. \quad (1.3)$$

Soon it was realized that the strategy-based definition of best response is epistemically inadequate: In many important types of games, such as one-shot non-cooperative games, players do not respond to other players' actual strategy choices, but rather to their *beliefs* about opponents' strategy choices⁷. This simple observation facilitated the emergence of epistemic game theory⁸, which then prompted the development of the epistemic definitions of basic concepts of orthodox game theory. The modern epistemic interpretation of a best response is that it is a strategy which maximizes player's expected utility, given player's consistent probabilistic beliefs about the opponents' strategy choices.

The epistemic concept of best response can be formally represented with an *epistemic model* of the game: A formal representation of each player's beliefs about the game and its players. There are two types of epistemic models that can be used to represent this idea – the state-space models and the type-space models. A state-space model will be used as a first example due to its relative prominence and intuitiveness.

⁷In non-cooperative one-shot games, players choose their strategies simultaneously, and so cannot observe each other's strategy choices before making their own strategy choices. In dynamic games, players cannot observe the future choices of their opponents, and, in some cases, even opponents' past strategy choices. When opponents' strategy choices cannot be observed, a rational player chooses a strategy which is optimal in light of his or her beliefs about the opponents' strategy choices. For extensive discussion, see Kreps 1990 and Fudenberg and Tirole 1991.

⁸Epistemic game theory is a subfield of game theory which focuses on the formal representation and analysis of players' belief structures and reasoning processes. One of the major objectives of epistemic game theory is to provide epistemic characterizations of the standard concepts of game theory, such as the Nash equilibrium. In epistemic game theory, non-cooperative games are treated as one-person decision problems under uncertainty. The combinations of opponents' strategies are viewed as possible states of the world, and the player is uncertain about the state of the world in which s/he has to make a strategy choice. The player is assumed to assign an internally consistent subjective probability distribution over the possible states of the world, and choose a strategy which maximizes his or her expected utility. For an extensive discussion of the principles and conceptual developments of epistemic game theory, see Dekel and Siniscalchi (2014) and Perea (2012). For a historical overview of the development of epistemic approach to game theory, see Perea (2014).

Let $\mathcal{M}^\Gamma(\Omega, \{\pi_i, \mathcal{H}_i, \zeta_i\}_{i \in I})$ be the state-space epistemic model of a normal form game Γ . The epistemic model consists of a set of possible states of the world Ω and, for every player $i \in I$, the information partition \mathcal{H}_i , a probability measure π_i on Ω , and a strategy function $\zeta_i : \Omega \rightarrow A_i$, where A_i is the set of actions of $i \in I$. Each state of the world $\omega \in \Omega$ includes a complete specification of all the parameters which may be the object of uncertainty on the part of any player of the game. Most importantly, each state ω includes a complete specification of actions chosen by every player of the game in that state. The information partition \mathcal{H}_i plays an important role in the characterization of player i 's beliefs and the state of the world: It assigns a set $h_i(\omega)$ to each $\omega \in \Omega$ in such a way that $\omega \in h_i(\omega)$ for all ω . The set $h_i(\omega)$ consists of those states of the world that i deems possible when the actual state of the world is ω . That is, if the true state of the world is $\omega \in h_i(\omega)$, then player $i \in I$ knows that the true state of the world is some element of $h_i(\omega)$, but s/he does not know which one it is⁹. The probability measure π_i is player i 's prior on Ω . Note that in the epistemic model players' strategies are treated as \mathcal{H}_i -measurable maps from states of the world to actions¹⁰.

Strategy ζ_i of player $i \in I$ is a best response if and only if, for every strategy $\tilde{\zeta}_i$ of player $i \in I$,

$$\sum_{\omega \in \Omega} \pi_i(\{\omega\}) u_i(\zeta_i(\omega), \zeta_{-i}(\omega)) \geq \sum_{\omega \in \Omega} \pi_i(\{\omega\}) u_i(\tilde{\zeta}_i(\omega), \zeta_{-i}(\omega)). \quad (1.4)$$

The epistemic concept of best response can also be characterized using the type-space epistemic model. In this model, it is assumed that each player holds not only beliefs about other players' strategy choices, but also beliefs about other players' beliefs about other players' strategy choices, beliefs about other players' beliefs about other players' beliefs, and so on *ad infinitum*. In other words, each player $i \in I$ is modelled as a decision-maker with an infinite hierarchy of beliefs about opponents' strategy choices and beliefs.

Formally, a type-space epistemic model of Γ can be defined as a tuple $\mathcal{E}^\Gamma = (\{T_i\}_{i \in I}, \{b_i\}_{i \in I})$, where T_i is the set of types of player $i \in I$ and $b_i : T_i \rightarrow \Delta(S_{-i} \times T_{-i})$ is a function which assigns, to every type $t_i \in T_i$, a probability measure on the set of opponents' strategy-type combinations. A set of opponents' strategy-type combinations is a cartesian product of a set

⁹For extensive technical discussion of the state-space model, see Aumann 1974, 1976 and 1987

¹⁰For an in-depth technical discussion of epistemic best response and subjective correlated equilibrium concepts, see Brandenburger and Dekel 1987.

$S_{-i} = \times_{j \in I \setminus \{i\}} S_j$ of possible combinations of opponents' strategies and a set $T_{-i} = \times_{j \in I \setminus \{i\}} T_j$ of possible combinations of opponents' types. The interpretation is that $b_i(t_i)$ represents a belief that type $t_i \in T_i$ holds about opponents' choices and beliefs¹¹. For each strategy $s_i \in S_i$, the expected utility of type $t_i \in T_i$, given his or her beliefs about the opponents' strategy combination, is

$$u_i(s_i, t_i) = \sum_{s_{-i} \in S_{-i}} (b_i(t_i))(s_{-i}) u_i(s_i, s_{-i}). \quad (1.5)$$

From definition (1.5) it follows that, given the belief about opponents' strategy choices of the type $t_i \in T_i$, strategy $s_i \in S_i$ is a best response to type t_i if and only if, for every strategy $\tilde{s}_i \in S_i$,

$$\sum_{s_{-i} \in S_{-i}} (b_i(t_i))(s_{-i}) u_i(s_i, s_{-i}) \geq \sum_{s_{-i} \in S_{-i}} (b_i(t_i))(s_{-i}) u_i(\tilde{s}_i, s_{-i}). \quad (1.6)$$

1.1.5 Best Response and The Nash equilibrium

Best-response reasoning plays a central role in the characterization of the Nash equilibrium – the central solution concept of game theory¹². The Nash equilibrium can be also defined both in terms of players' strategies and in terms of players' beliefs. In terms of strategies, the Nash equilibrium is a strategy profile where each player's strategy is a best response to a combination of other players' strategies. In other words, the Nash equilibrium is a strategy profile such that no player has an incentive to unilaterally deviate by changing his or her strategy. A mixed strategy profile $\sigma^* = (\sigma_1^*, \dots, \sigma_m^*)$ is a Nash equilibrium of Γ if and only if, for every player $i \in I$, $u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(s_i, \sigma_{-i}^*)$ for all $s_i \in S_i$ ¹³.

The aforementioned characterization of the Nash equilibrium, due to Nash (1950b, 1951), is now sometimes referred to as an *ex post* definition of the Nash equilibrium, since each player's strategy is defined as a best response to a fixed combination of opponents' actual strategies. Some critics have argued that this definition implies that each player somehow knows the opponents' strategy choices before those choices are actually made, and that

¹¹For an in-depth technical discussion of the type-space epistemic models, see Perea 2012 and Dekel and Siniscalchi 2014.

¹²A Nash equilibrium refinement is a Nash equilibrium which has a specific set of mathematical properties, and so the concept of best-response reasoning is central to all of the Nash equilibrium refinements. For extensive discussion of equilibrium refinements, see Myerson 1991 and Weibull 1995.

¹³A pure strategy equilibrium is a profile of pure strategies which satisfies the same condition. For extensive discussion, see Myerson 1991.

it is difficult to incorporate this epistemic assumption into a game theoretic analysis of those games where players choose their strategies simultaneously and without communication (see, for example, Perea 2014).

In a more recent epistemic characterization of the Nash equilibrium, this solution concept is defined in terms of players' *conjectures* – subjective probability distributions over opponents' strategy choices. Aumann and Brandenburger have shown that if players' payoff functions, rationality and their conjectures are *mutually known* in a two player game, then players' conjectures constitute a Nash equilibrium. In a game with more than two players, players' conjectures always constitute a Nash equilibrium if payoff functions and rationality of the players are mutually known, each player's conjecture is *commonly known*, and players' have a common prior (for detailed discussion and proofs, see Aumann and Brandenburger 1995). This result allows for an *ex ante* characterization of the Nash equilibrium, according to which *Nash equilibrium* is a combination of players' beliefs. Aumann and Brandenburger's seminal work allows us to talk about the players of the game being in equilibrium *before* their strategy choices are observed¹⁴.

The Nash equilibrium concept has several important functions in both normative and descriptive game theory. In formal game theoretic analysis, the Nash equilibrium concept is the definition of the rational solution of a game: If a game has a solution that is common knowledge among the rational players, it must necessarily be a Nash equilibrium (Binmore 2005). The logic behind this requirement is fairly obvious: If the strategy profile of the game known to the players is not a Nash equilibrium, at least one player will have an incentive to deviate, which means that non-equilibrium strategy profile, by definition, cannot be the solution of the game. In fact, an even stronger argument for the Nash equilibrium concept, supported by some game theorists, is that rational players who know about each other's rationality should only expect the outcome of the game to be a Nash equilibrium, at least in cases where rationality is common knowledge. In other words, Nash equilibrium is assumed to be a consequence of common knowledge of rationality. Other theorists, however, reject this assumption as being

¹⁴For an extensive discussion of the epistemic interpretation of the Nash equilibrium, see, for example, Perea 2012.

too strong¹⁵.

Despite the differing opinions on whether players' choice of strategies constituting the Nash equilibrium can be viewed as a consequence of rationality, there is a wide consensus among the game theory experts that Nash equilibrium plays a fundamental role as a *stability concept* which, in descriptive models of game theory, explains the recurring patterns of social agents' behaviour. The underlying idea is that the reason of why social agents repeatedly choose the same behavioural response to a certain type of interdependent decision problem is that social agents' responses constitute a Nash equilibrium – a combination of individuals' actions, such that no interacting individual can find a personally more advantageous action, given the expected actions of other individuals¹⁶.

Over the last couple of decades, several important conceptual relations between the Nash equilibrium and the stability concepts of evolutionary game theory have been established. In evolutionary game theory, the concept of evolutionarily stable state is used to characterize the dynamic properties of the population. A population is in an evolutionarily stable state if, after a disturbance, it returns to playing a certain strategy or mix of strategies, provided that the disturbance is not too large (for a technical discussion, see Thomas 1984 and Weibull 1995). A closely related concept to evolutionar-

¹⁵For example, Bernheim (1984) argues that Schelling's (1960) work on focal points involves an assumption that players share a common belief that the outcome of the game must be a Nash equilibrium. An assumption that players only consider the Nash equilibria of the game can also be found in Harsanyi and Selten's work on equilibrium selection principles (1988). It is important to note that the idea that rational players should expect the outcome of the game to be a Nash equilibrium has not been endorsed by the proponents of epistemic game theory, some of whom have argued that Nash behaviour is neither a necessary consequence of rationality, nor even a plausible empirical proposition. For an extensive discussion of the relationship between common knowledge of rationality and the Nash equilibrium, see, for example, Ellsberg 1956, Luce and Raiffa 1957, Bernheim 1984, and Perea 2014.

¹⁶This explanatory strategy can be applied to both one-shot and repeated games. For example, Binmore (2005) applies this explanatory strategy to repeated games in order to explain the emergence of social norms. Bicchieri (2006) and Gintis (2008) offer game theoretic accounts of social conventions based on one-shot games. Fehr and Schmidt's (1999) inequity aversion model shows that a relatively frequently observed cooperative outcome of the experimental Prisoner's Dilemma game (a game where, for both players, cooperation is irrational in the orthodox game theoretic sense) can be modelled as a Nash equilibrium of a game, the payoff structure of which captures players' motivation to avoid unequal distribution of material (e.g. monetary) payoffs – a motivation which is not represented in the payoff structure of the original experimental Prisoner's Dilemma game (for details, see Fehr and Schmidt 1999). In all cases, the persistence of a certain pattern of behaviour is explained as a consequence of players' actions constituting a Nash equilibrium of the game.

ily stable state is evolutionarily stable strategy – a strategy which, if used by every individual in the population, makes the population uninvadable, over evolutionary time, by mutants using other strategies, provided that the initial proportion of invading mutants is not too large (see, for example, Maynard Smith and Price 1973)¹⁷. A population state where everyone plays an evolutionarily stable strategy is an evolutionarily stable state. It has been shown that a set of evolutionarily stable states is a subset of the Nash equilibria of the evolutionary game. Arguably the most important result is that any state where everyone plays an evolutionarily stable strategy is a perfect and proper Nash equilibrium of the evolutionary game, and that every strict Nash equilibrium¹⁸ of the evolutionary game is an evolutionarily stable state. Therefore, if everyone in the population plays a strict Nash equilibrium strategy, the population is in an evolutionarily stable state (for discussion and proofs, see van Damme 1987 and Weibull 1995). From the aforementioned theoretical results it follows that, since the evolutionary process leads the population to converge to an evolutionarily stable state, the population in the evolutionarily stable state must be playing a Nash equilibrium of the game.

Evolutionary game theory is now widely viewed as providing an important non-rationalistic justification of the Nash equilibrium concept: Even if players were strategically unsophisticated agents (in fact, even if they were unable to reason at all), the evolutionary process should, over time, lead the population to an evolutionarily stable state – a state which is a Nash equilibrium of the population game. A considerable number of social phenomena, such as social norms and conventions, are viewed as products of a lengthy process of cultural and possibly even biological adaptation (for example, see Gintis 2008). Social agents who repeatedly respond to a certain type of interdependent decision problem in the same way are viewed as playing a Nash equilibrium of a certain recurring game¹⁹. Any theory suggesting that a decision-making mechanism which leads individuals to an out-of-equilibrium behaviour could have been selected due to its evolution-

¹⁷A population state which is asymptotically stable is an evolutionarily stable state. It has been shown that every regular evolutionarily stable strategy is asymptotically stable under the replicator dynamic (Taylor and Jonker 1978, Hines 1980), general imitative dynamic (Cressman 1997), any impartial pairwise comparison dynamic, such as the Smith dynamic, any separable excess payoff dynamic, such as the Brown-von Neumann-Nash dynamic, and under the best response dynamic (Sandholm 2010b).

¹⁸A strategy profile $\mathbf{s}^* = (s_1^*, \dots, s_m^*)$ is a strict Nash equilibrium if and only if, for every $i \in I$, it is the case that $u_i(s_i^*, s_{-i}^*) > u_i(s_i, s_{-i}^*)$ for all $s_i \in S_i$.

¹⁹For a critical overview of the use of evolutionary game theoretic models in explanations of social phenomena, see, for example, Alexander 2007.

ary advantages is therefore generally viewed as being deeply conceptually problematic²⁰.

Evolutionary justification aside, any theory which claims to show that player's decision to *knowingly*²¹ play a part in an out-of-equilibrium strategy profile is a reasonable response to certain types of interdependent decision problems is viewed as objectionable because of another reason – its incompatibility with the orthodox conception of rationality.

1.1.6 Rationality as Best-Response Reasoning

In orthodox game theory, a rational player is a best-response reasoner. Best-response reasoning is a process by which a rational player arrives at his or her strategy choice. Best-response reasoner is a player who always chooses a strategy which is a best response – a strategy which maximizes player's expected utility, given his or her consistent probabilistic beliefs about opponents' strategy choices.

In most rationalistic game theoretic models, players' rationality is assumed to be common knowledge or, in more recent epistemic models, a *common belief*. That is, each player is assumed to believe that every other player of the game is rational, believe that every other player believes that every other player of the game is rational, and so on *ad infinitum* (for an in-depth technical discussion of the role the common belief in rationality assumption in game theoretic analysis, see Perea 2012). Common knowledge or a common belief in rationality assumption implies that each player

²⁰The idea that a mode of reasoning leading the individuals to out-of-equilibrium outcomes could have been selected for its group fitness enhancing property has been defended by, among others, Gauthier (1986), McClennen (1988) and Bacharach (2006). The argument rests on the observation that certain mixed motive games, such as the Prisoner's Dilemma game, have a unique Pareto inefficient Nash equilibrium and a Pareto-optimal out-of-equilibrium outcome which is mutually beneficial – each player gets a higher expected payoff than the one associated with a Pareto inefficient Nash equilibrium of the game. However, the out-of-equilibrium outcome creates an incentive for each player to unilaterally deviate in order to maximize the personal expected payoff, at the expense of the player who resists the temptation to deviate. The incentive to deviate thus makes cooperation impossible. A mechanism which allows all the players of the game to resist, for whatever reason, the temptation to maximize their personal expected utility and cooperate in implementing a mutually beneficial outcome increases the fitness of the group, making it more likely to succeed at the group selection level. This idea will be discussed in considerable detail in chapter 2.

²¹A substantial number of game theorists reject the idea that an outcome of the game played by rational players must necessarily be a Nash equilibrium. The reason of this is that players' conjectures (i.e. beliefs about opponents' strategy choices) are, in many cases, private information. This possibility will be discussed in subsection 1.1.6.

knows/believes that every player of the game is a best-response reasoner.

Common belief in rationality is a weaker epistemic assumption than common knowledge of rationality. The common knowledge assumption implies that the player cannot be mistaken in his or her belief that every opponent is rational, while a probability 1 belief may be either true or false. In recent years, the common belief of rationality assumption tends to be preferred over the common knowledge of rationality assumption due to its purported realism. The conceptual differences between the common knowledge of rationality and the common belief in rationality concepts will not be important in the following discussion, and so these two concepts will be used as interchangeable terms²².

The aforementioned definition of best response reasoning implies that a best-response reasoner never chooses a strategy which is not a best response. It also implies that a best-response reasoner never chooses a non-rationalizable strategy – a strategy which is never optimal, irrespective of what probabilistic beliefs the player holds about opponents’ strategy choices (for technical definition, see Bernheim 1984 and Pearce 1984).

In the standard analysis of non-cooperative games, each player’s probabilistic beliefs about the opponents’ strategy choices, or conjectures, are assumed to be private information. That is, players are assumed not to know each other’s conjectures. However, common knowledge of rationality allows the players to eliminate non-rationalizable strategies from strategic considerations. Common knowledge of rationality implies that every player knows that none of the opponents will choose a non-rationalizable strategy, knows that every opponent knows this, and so on *ad infinitum*. Common knowledge of rationality assumption is central to the iterative elimination of strictly dominated strategies, which is a procedure used to identify the set of rationalizable strategies of the game – strategies, the choice of which is consistent with a common knowledge of rationality assumption²³.

Best-response reasoning assumption is one of the reasons of why a commonly known rational solution of the game in the orthodox game theoretic sense must be a Nash equilibrium: A non-equilibrium solution of the game creates an incentive for at least one of the rational players to maximize his

²²For an extensive technical discussion of the common belief concept, see Samet 2013.

²³In two player games, iterated elimination of non-rationalizable strategies coincides with the iterated elimination of strictly dominated strategies. This relationship, however, does not hold in games with more than two players. Every strategy which gets eliminated in the process is non-rationalizable, but the converse is not necessarily true: A strategy may survive the iterated elimination procedure, and yet be non-rationalizable. For a detailed discussion, see Bernheim 1984, Pearce 1984, Fudenberg and Tirole 1991 and Perea 2012.

or her expected utility by choosing a different strategy, which means that the commonly known solution of the game will not be implemented via players' joint actions. A player who knows that the opponent is rational should expect him or her to deviate in a situation where a unilateral deviation is beneficial, and respond by choosing an optimal response to the expected deviation. In other words, the commonly known out-of-equilibrium solution of the game is inherently *unstable* due to being incompatible with the individual incentives of best-response reasoners (for extensive discussion, see Myerson 1991 and Bicchieri 1995).

The question of whether a Nash equilibrium is a necessary outcome of a complete information game played by rational players, and whether rational players should expect a Nash equilibrium to be the outcome of such a game, requires a more thorough examination. The currently prevailing position is that common knowledge of rationality and of the payoff structure of the game does not imply that a Nash equilibrium is going to be played. To see why this is the case, consider the following two player Weak Dominance game depicted in Figure 1.2:²⁴:

	<i>t1</i>	<i>t2</i>
<i>s1</i>	10, 0	5, 2
<i>s2</i>	10, 1	2, 0

Figure 1.2: Weak Dominance game

In this game, the row player has to choose between pure strategies *s1* and *s2*, while the column player has a choice between pure strategies *t1* and *t2*. The game has two pure strategy Nash equilibria (*s1, t2*) and (*s2, t1*) and a mixed strategy Nash equilibrium $(\frac{1}{3}s1, \frac{2}{3}s2; t1)$. There are no strictly dominated strategies²⁵, which means that common knowledge of rationality alone does not allow the players to eliminate any of the strategies from their strategic considerations. Suppose that the row player believes that the

²⁴This game is due to Kreps 1990.

²⁵Strategy *s2* is weakly dominated for the row player: It gives the same payoff against strategy *t1* as strategy *s1*, but a strictly worse payoff against strategy *t2*. However, strategy *s2* cannot be eliminated on the basis of common knowledge of rationality alone, since it is a best response against a conjecture that the column player will play strategy *t1* with probability 1. For an in-depth discussion of the epistemic assumptions underlying the elimination of weakly dominated strategies, see Perea 2012.

probability of the column player choosing strategy $t1$ is 1²⁶. The row player should then be indifferent between playing strategies $s1$ and $s2$, since the expected payoff from playing either of the two strategies would be the same – 10. Suppose that the column player believes that the probability of the row player choosing $s2$ is 1. In that case, the column player’s best response is $t1$. If the row player plays $s1$, the players end up playing the strategy profile $(s1, t1)$ which, unlike the strategy profile $(s2, t2)$, is not a Nash equilibrium of the game. However, the players could not be blamed for being irrational, since each player’s strategy is optimal in light of his or her private belief about the other player’s strategy choice²⁷.

1.2 Social Coordination and the Conceptual Limitations of the Best-Response Reasoning Models

The aforementioned example indicates a more general issue with the standard game theoretic analysis based on best-response reasoning: In every game with multiple rationalizable outcomes, standard game theoretic analysis is indeterminate. More specifically, the theory cannot predict, which combination of strategies will be the outcome of a gameplay. The reason of this result can be traced to the logic of best-response reasoning. Each best-response reasoner is supposed to choose an optimal action in light of his or her beliefs about the opponents’ strategy choices. The problem is that each player’s choice depends on his or her beliefs about each opponent’s strategy choice, and each opponent’s strategy choice depends on his or her beliefs about the player’s own strategy choice. From the common knowledge of rationality and of the payoff structure of the game, the player cannot infer the opponents’ strategy choices, and so s/he is left in a state of strategic uncertainty having no indication as to what the opponents can be expected to do. Consequently, the player cannot decide, which one of the available strategies is a best response (for extensive discussion of this problem, see Bicchieri 1995 and Bacharach 2006).

In more recent epistemic game theory models, the players are assumed

²⁶It is important to note that common knowledge of rationality implies that each player assigns zero probability to the event of the opponent choosing a strictly dominated strategy. It does not, however, imply that each player must assign a strictly positive probability to every rationalizable strategy. For extensive discussion, see Perea 2012.

²⁷For a technical discussion of the difference between *ex ante* and *ex post* optimality of choice, see, for example, Brandenburger and Dekel 1987 and Perea 2012.

to have a prior on the set of possible states of the world (in type space model, on the set of opponents' strategy-type combinations). A prior does not automatically imply that players will end up playing a Nash equilibrium of the game, or that they will successfully coordinate their actions. The players may have different priors and choose strategies which, although *ex ante* optimal, do not constitute a Nash equilibrium. In some cases, the players may end up with an undesirable outcome. To see why this might be the case, consider the two player coordination game depicted in Figure 1.3:

	<i>t1</i>	<i>t2</i>
<i>s1</i>	2, 1	1, 2
<i>s2</i>	4, 3	0, 1

Figure 1.3: Coordination game

In this game, the row player has to choose either to play strategy *s1* or strategy *s2*, while the column player has to choose between strategy *t1* and strategy *t2*. This game has two Nash equilibria in pure strategies – (*s1*, *t2*) and (*s2*, *t1*)²⁸.

Suppose that, for each player, the set Ω contains two possible states of the world – ω_1 and ω_2 . In state $\omega_1 \in \Omega$ the opponent plays strategy *s1/t1* and in state $\omega_2 \in \Omega$ the opponent plays strategy *s2/t2*. Suppose that row player's prior assigns probability 3/4 to state ω_1 and probability 1/4 to state ω_2 , while column player's prior assigns probability 5/6 to state $\omega_1 \in \Omega$ and probability 1/6 to state $\omega_2 \in \Omega$. In such case, row player's best response is strategy *s2*, column player's – *t2*. Strategy profile (*s2*, *t2*) is not a Nash equilibrium. It is also the worst possible outcome for both players in this game. However, each player's response is *ex ante* optimal.

With certain priors, coordination is, of course, possible. In epistemic models of game theory, a *common prior* is a fairly standard assumption. Recall that in state-space epistemic models, each player $i \in I$ is assumed to have a prior probability distribution p_i over the set of possible states of the world Ω (in type-space models, over the set of opponents' strategy-type combinations). An information structure is said to have a common prior if

²⁸The game also has a mixed Nash equilibrium ($\frac{2}{3}s1, \frac{1}{3}s2; \frac{1}{3}t1, \frac{2}{3}t2$). However, in epistemic models of one-shot non-cooperative games, it is standard to treat mixed strategies not as players' actual randomizations over their pure strategies, but as players' probabilistic beliefs about opponents' actions – subjective probability distributions over opponents' pure strategies. In other words, players are assumed to use only pure strategies in one-shot games. For extensive discussion, see Osborne and Rubinstein 1994 and Perea 2012.

$p_i = p$ for all $i \in I$. With common prior assumption, coordination can be achieved. For example, if both players were to assign probability 1/2 to both state ω_1 and state ω_2 , the row player's best response would be strategy s_2 while the column player's best response would be strategy t_1 . The players would thus end up playing strategy profile (s_2, t_1) – a Pareto efficient²⁹ Nash equilibrium of this game (for extensive discussion, see Aumann 1987 and Brandenburger and Dekel 1987).

A common prior assumption imposes substantial restrictions on players' possible beliefs. Aumann (1976) has shown that if two players are Bayesian-rational decision-makers and have a common prior, and if their posteriors are common knowledge, then their posteriors must be equal. In other words, with a common prior, it cannot be common knowledge that different players hold different posterior beliefs about any event (for discussion and proofs, see Aumann 1976).

A common prior assumption is based on Harsanyi's doctrine: An idea that any two rational players who have access to the same information will necessarily have the same probabilistic assessment of the situation, and any difference in players' assessments must necessarily be due to differences in the information available to them (for extensive discussion, see Harsanyi 1967, 1968a,b, Morris 1995 and Bicchieri 1995). The problem with the common prior assumption is that it offers very little in terms of explaining how people coordinate their actions. Common knowledge of rationality and of the payoff structure of the game does not provide enough information for the players to form beliefs about each other's strategy choices. It is therefore not clear, what kind of information about the non-cooperative game which is available to the players justifies the common prior assumption, let alone the common prior which leads the players to choosing strategies constituting a Nash equilibrium. As has been pointed out by Bicchieri (1995), even in situations where it is common knowledge that players only consider the set of Nash equilibria as possible solutions, the question of how they should assign subjective probabilities to opponents' actions in games with multiple Nash equilibria does not have a clear answer (for details, see Bicchieri 1995). In other words, an assumption that players have a common prior which allows them to coordinate their actions is a substantial assumption which requires a compelling conceptual justification – a justification which the theory of

²⁹An outcome of a game is Pareto efficient, or Pareto optimal, if there is no other outcome available which would make some player better-off without making anyone else worse-off.

games itself does not offer³⁰.

Because of the conceptual limitations of the best-response reasoning model, the theory is sometimes criticized for its inability to single out unique solutions in games with multiple Nash equilibria, even in those which, at least intuitively, seem to have an ‘obvious’ solution. A canonical example of a game where the standard game theoretic analysis leads to conclusions which contradict our common sense intuitions concerning rational behaviour is the Hi-Lo game depicted in Figure 1.4:

	<i>hi</i>	<i>lo</i>
<i>hi</i>	2, 2	0, 0
<i>lo</i>	0, 0	1, 1

Figure 1.4: Hi-Lo game

In this game, two players must simultaneously and independently choose either to play *hi* or to play *lo*. If both players choose *hi*, each gets a payoff of 2. If both choose *lo*, each gets a payoff of 1. If one player chooses *hi* while the other chooses *lo*, both players get a payoff of 0 – the worst payoff attainable in this game.

Both *hi* and *lo* are rationalizable strategies. The game has two pure strategy Nash equilibria – (*hi*, *hi*) and (*lo*, *lo*). The third one is a mixed strategy Nash equilibrium, in which players randomize between pure strategies *hi* and *lo* with probabilities 1/3 and 2/3 respectively. The mixed Nash equilibrium yields each player an expected payoff of 2/3. From the perspective of orthodox game theory, every Nash equilibrium is a rational solution of this game: It is rational for a player to play *hi* with probability 1/3 and *lo* with

³⁰One of the attempts to justify the common prior assumption is based on idea that individuals coming from the same population may have pre-formed beliefs about the behaviour of the members of the population in situations of a certain type. In other words, individuals sharing similar social experiences will likely have similar expectations concerning each other’s behaviour, and therefore may successfully coordinate their actions (see, for example, Bicchieri 2006 and Gintis 2008). The problem with this justification of the common prior assumption is that it rests on empirical claim about the population state – a claim which may or may not be true and which may be difficult, if not impossible, to test. This justification of the common prior assumption will be discussed in considerable detail in chapter 4.

probability $2/3$ if s/he believes that the opponent will do the same thing³¹. It is also rational to play *lo* if s/he believes that the probability of opponent playing *lo* is higher than $2/3$. For many people, however, pure strategy Nash equilibrium (lo, lo) and the mixed strategy Nash equilibrium do not appear as convincing solutions of what they perceive to be a *common interest* game: Nash equilibrium (hi, hi) is clearly the best outcome for both players, and there is no conflict of players' interests in this game. The perfect alignment of player's personal interests seems to be the primary reason of why an idea that a rational player could expect another rational player to choose *lo* is not compelling.

The critics of the orthodox game theory argue that the failure of game theory to agree with our 'high quality intuitions' concerning rational behaviour should be viewed as a conceptual limitation of the theory, which should not be dismissed as insignificant and, ideally, addressed with a model of reasoning compatible with our intuitions about rational behaviour in non-cooperative games (for a defense of this view, see, for example, Olcina and Urbano 1994 and Bacharach 2006).

The aforementioned critique, however, should be taken with a grain of salt: It relies on the interpretation of game theory as aiming to provide a 'unique rational recommendation on how to play' (Bicchieri 1995: 316). Strictly speaking, best-response reasoning, which is the main target of criticism, is not a model of strategic reasoning, but rather a consistency requirement. A rational player is supposed to have internally consistent probabilistic beliefs about opponents' strategy choices, and never choose a strategy which is not optimal in light of those beliefs. A rational player who knows the opponents to be rational is only supposed never to expect them to play non-rationalizable strategies (Bicchieri 1995). Notice, however, that best-response reasoning model is not a theory of how rational players form beliefs about opponents' choices of rationalizable strategies (Olcina and Urbano 1994). Therefore, it is not surprising that the theory cannot single out unique solutions in games with multiple Nash equilibria, or that it cannot address a criticism that certain Nash equilibria are not compelling solutions due to intuitively unreasonable beliefs that the players would have to hold in order to end up playing them. It can be argued that the question of what a rational player should believe about opponents' rationalizable strategies

³¹Strictly speaking, if the player believes that the opponent will randomize between *hi* and *lo* with probabilities $1/3$ and $2/3$ respectively, the player is indifferent between playing *hi* or *lo* with probability 1 and playing strategy *hi* with probability $1/3$ and *lo* with probability $2/3$. The expected payoff associated with each option is the same – $2/3$. This result raises the question of whether it is reasonable to assume that players would play randomized strategies in a one-shot game.

falls *outside* of the scope of game theory or, even if it does not, that there is no uniquely rational belief-formation model for non-cooperative games (for a defence of a latter position, see, for example, Perea 2012).

However, the criticism is valid in a sense that the question of how real-world decision-makers choose among rationalizable strategies is of fundamental importance for the descriptive branch of game theory. Experiments suggest that people are fairly successful in coordinating their actions in games with multiple rationalizable outcomes³². The aim of the descriptive models of game theory is to explain people's behaviour, and so they should provide definitive theoretical predictions concerning people's actions. Therefore, the question of how social agents manage to coordinate their actions in non-cooperative games is an important explanatory challenge for the descriptive branch of game theory.

1.3 Social Coordination Theories

Currently there are several descriptive theories which purport to explain how people resolve games with multiple rationalizable outcomes. These theories will be briefly introduced in this section. This introduction should not be viewed as a comprehensive overview: Most of the theories have multiple different versions, and this overview will cover only the general principles of each approach.

1.3.1 Cognitive Hierarchy Theory

The cognitive hierarchy theory³³ suggests that players can be classified into types on the basis of their ability to engage in strategic reasoning. The theory postulates a hierarchy of cognitive levels (0, 1, 2...). Each type of player can be characterized by a cognitive level, which represents the degree to which that type of player can reason about the other players. Level 0 players are unable to engage in strategic reasoning at all. They are assumed to be randomizing between strategies according to some exogenous probability distribution p_0 . The standard assumption is that probability distribution p_0

³²For experimental results suggesting that people can achieve high rates of coordination in non-cooperative games with multiple Nash equilibria, see, for example, Mehta et al. 1994, Bacharach and Bernasconi 1997, Colman and Stirk 1998, Crawford et al. 2008, Bardsley et al. 2010, Cabon-Dhersin and Etchart-Vincent 2010, Faillo et al. 2013, 2016.

³³For early and more recent theoretical developments and experimental tests of cognitive hierarchy theory, see Nagel 1995, Stahl and Wilson 1994, 1995, Ho et al. 1998, Costa-Gomes et al. 2001, Camerer et al. 2004, Crawford et al. 2008, Bardsley et al. 2010, Faillo et al. 2013, 2016.

is uniform, which means that level 0 players choose each pure strategy with equal probability (see, for example, Stahl and Wilson 1995 and Camerer et al. 2004). It is often assumed that distribution p_0 is common knowledge, although in some versions of the theory this assumption is replaced with a presumably more realistic assumption that players may have different beliefs about the distribution p_0 (see, for example, Bardsley et al. 2010). In some versions of the theory, a *non-uniform* probability distribution is used to represent a tendency of unsophisticated players to opt for strategies that they perceive as being salient for some non-strategic (e.g. psychological) reason (see, for example, Crawford et al. 2008).

A level 1 player believes that all the opponents are level 0 reasoners who randomize between pure strategies according to p_0 . Level 1 player chooses a strategy which, given this belief, maximizes his or her expected payoff.

A level 2 player believes that each opponent is either level 0 or level 1 player. With probability $q_0/(q_0 + q_1)$ the opponent is a level 0 player who randomizes between pure strategies according to distribution p_0 (it is important to note that this theory relies on assumption that $q_0 > 0$). With probability $q_1/(q_0 + q_1)$ the opponent is a level 1 player who expects the opponents to be level 0 players and chooses a best response to that belief. Level 2 player chooses a strategy which, given his or her probabilistic beliefs about opponents' types, maximizes his or her expected payoff.

A level 3 player believes that opponents can be either level 0, level 1 or level 2 players and, like level 1 and level 2 players, chooses a best response to his or her probabilistic beliefs about the opponents' types. The same principles of modelling apply to reasoning of players of higher cognitive levels³⁴.

Cognitive hierarchy theory rests on a fundamental assumption that there is no type of player who can conceive the possibility that other players may be either of the same or of the higher cognitive level that s/he is. In other words, each possible type of player is boundedly rational in a sense that s/he can only think of opponents as being less strategically sophisticated than s/he is (Camerer et al. 2004). Another important assumption is that level $k \geq 1$ players 'have an accurate guess about the relative proportions of players who are doing less thinking than they are' (Camerer et al. 2004:

³⁴In a simpler version of the theory, it is assumed that level 1 player best-responds to beliefs about the behaviour of level 0 players, level 2 player best-responds to beliefs about the behaviour of level 1 player, and so on. In other words, each $k \geq 1$ level player is assumed to believe that all the other players are level $(k - 1)$ players, and chooses a best response accordingly. This simplifies the model in a sense that players are not uncertain about opponents' types. In some versions of the theory, it is assumed that no player can be of higher cognitive level than level 2. See, for example, Crawford et al. 2008.

864).

The simplest version of the cognitive hierarchy theory – the one which assumes that probability distribution p_0 is uniform – can explain coordination in certain types of coordination games, such as in common interest games with a unique Pareto optimal Nash equilibrium. For example, it explains coordination in the aforementioned Hi-Lo game (Figure 1.4). For level 1 player, the optimal response in the Hi-Lo game is to play strategy *hi*, since s/he expects the level 0 opponent to play *hi* with probability 1/2. A level 2 player expects the opponent to be either level 0 or level 1 player. Level 2 player expects the level 1 player to play *hi*, and so his or her optimal response is *hi*. Just like level 1 player, level 2 player expects the level 0 opponent to play strategy *hi*, and so his or her best response is, again, *hi*. Strategy *hi* is therefore level 2 player’s best response to both level 1 and level 0 player’s actions. It can easily be checked that strategy *hi* is a best response for a player of any higher cognitive level, irrespective of what probability distribution s/he assigns over the opponent’s cognitive types (given the constraint that $q_0 > 0$).

A probability distribution p_0 which represents a tendency of unsophisticated players to choose salient strategies more frequently than other available strategies can be used to account for the role of non-payoff-relevant factors, such as strategy labels or their position in the game matrix³⁵.

One of the perceived advantages of the cognitive hierarchy theory is that it retains some of the elements of orthodox best-response reasoning model: Although players are assumed to have limited reasoning abilities and therefore cannot be treated as Bayesian-rational decision-makers, the players whose cognitive level is higher than 0 are assumed to make choices consistent with the expected utility maximization principle. In other words, the players are bounded best-response reasoners: They choose optimal responses to their beliefs about opponents’ actions, even though their ability to form consistent beliefs about the opponents’ types is restricted by their cognitive limitations.

The theory can be criticized on empirical as well as on conceptual grounds. Some critics have pointed out that the theory fails to explain coordination in games with multiple payoff-identical Pareto optimal outcomes and a unique Pareto suboptimal outcome. For example, experiments suggest that in a two player game where players have to coordinate by choosing the same payoff outcome from a set of possible outcomes $\{(10, 10), (10, 10), (9, 9)\}$,

³⁵For example, in a version of the theory suggested by Crawford, Gneezy, and Rottenschreich (2008), level 0 players are assumed to choose payoff salient strategies and use label salience to resolve cases where two or more strategies are payoff salient.

people often coordinate their actions on a Pareto suboptimal outcome $(9, 9)$. According to the principles of cognitive hierarchy theory, level 1 or higher cognitive level players would not choose outcome $(9, 9)$ (for extensive discussion of this criticism, see, for example, Bardsley et al. 2010 and Faillo et al. 2016).

The theory also faces difficulties explaining high coordination rates in games where players have conflicting preferences over the Pareto optimal outcomes. For example, consider the following two player three strategy coordination game depicted in Figure 1.5:

	l	c	r
u	11, 9	0, 0	0, 0
m	0, 0	10, 10	0, 0
d	0, 0	0, 0	9, 11

Figure 1.5: Coordination game

This game has three Pareto optimal pure strategy Nash equilibria – (u, l) , (m, c) and (d, r) . Experiments suggest that an absolute majority of people end up playing the Nash equilibrium (m, c) (see, for example, Faillo et al. 2013). The cognitive hierarchy theory cannot explain this result. Consider the game from the row player’s perspective. Level 1 row player should expect that level 0 column player will play each pure strategy with probability $1/3$ and respond by playing u . A level 2 row player should expect the column player to be either level 0 or level 1 player. If the row player believes that the column player is a level 0 player, his or her best response is strategy u . If the row player believes that the column player is a level 1 player, s/he expects the column player to choose strategy r , and so his or her best response is d . Since the row player is uncertain about the opponents’ cognitive level, s/he will choose either u or d , depending on the probability distribution over the column player’s cognitive types: The row player will choose strategy u if the probability of the opponent being a level 0 player is higher than $27/29$ and strategy d if the probability of the opponent being a level 0 player is lower than $27/29$. For level 3 row player, a best response will also be either strategy u or strategy d , depending on the probability distribution over column player’s types. It is easy to check that the same result holds for any player of higher cognitive level.

A level 1 column player should always choose strategy r . For a column player of any higher cognitive level, a best response can be either strategy l or strategy r , depending on the probability distribution over row player’s

cognitive types. Notice that level $k \geq 1$ row player never plays strategy m and level $k \geq 1$ column player never plays c . The theory is thus unable to explain why approximately 86% of people end up playing the Nash equilibrium (m, c) (for experimental results, see, for example, Faillo et al. 2013).

The most serious *conceptual* criticism of the theory can be directed against the suggested structure of players' beliefs about the opponents' types. In cognitive hierarchy theory, level $k \geq 1$ player's beliefs are anchored in beliefs about level 0 player's behaviour. The explanatory power of the theory rests on assumptions that the proportion of level 0 players in the population is sufficiently high. In some versions of the theory, it is assumed that level 0 players do not exist at all, and the population is a mixture of level 1 and level 2 players (see, for example, Crawford et al. 2008). This assumption implies that every level $k \geq 1$ player holds incorrect beliefs about the opponents' types. The theory offers no explanation of how players form a belief that a non-existing player type is present in the population, or why players do not realize, over time, that level 0 players do not exist. This version of the theory also implies that players are not only boundedly rational when it comes to reasoning about the opponents' cognitive abilities, but also incapable of checking their beliefs against the available information: If players were able to update their beliefs on information obtained by observing other players' behaviour, the proportion of players believing in the presence of level 0 players in the population would decrease over time. The cognitive hierarchy theory would become explanatory irrelevant in the long run.

In other versions of the theory, the possibility that level 0 players are actually present in the population is not ruled out. However, an idea that a non-negligible number of real-world players cannot engage in strategic reasoning and choose their strategies at random seems empirically implausible, particularly when used to explain players' behaviour in simple common interest games, such as the Hi-Lo game. Therefore, a level $k \geq 1$ player's belief that a non-negligible proportion of social agents are level 0 players must, in most cases, be false.

It is important to note that the criticism directed against a belief in the non-negligible presence of level 0 players has no relevance if cognitive hierarchy theory is interpreted merely as a model offering an approximately accurate description of observed behaviour. In that case, any worries about the realism of the assumptions characterizing players' beliefs can be dismissed as irrelevant. However, the criticism is valid if we interpret the theory as offering an approximate description of the process of reasoning by which different types of players arrive at their action choices.

1.3.2 Stackelberg Reasoning

Stackelberg reasoning, or Stackelberg heuristic, suggests that each player chooses a strategy on the basis of a belief that the opponent will correctly predict his or her strategy choice and choose a best response to it. If the player uses Stackelberg heuristic, s/he expects the opponent to always play the same Nash equilibrium as s/he does. Given such expectations, player's choice of a strategy which constitutes a Nash equilibrium associated with the highest personal payoff is an optimal response. For example, consider the Hi-Lo game (Figure 1.4). A player who uses Stackelberg heuristic should choose strategy *hi*, since s/he should expect the opponent to predict the choice of *hi* and respond by playing *hi* as well (for extensive discussion, see Colman and Bacharach 1997 and Colman 1997).

One of the perceived advantages of the theory is that it retains some of the basic principles of the orthodox best-response reasoning model: Each player who uses Stackelberg reasoning is assumed to be choosing a *Stackelberg strategy* – a strategy which maximizes his or her expected payoff in light of a belief that the opponent will correctly guess his or her strategy choice and choose a best response to it (Colman and Bacharach 1997). The theory, however, can be criticized on both empirical and conceptual grounds. The proponents of the theory suggest that decision-makers only use Stackelberg heuristic in the so called *Stackelberg-soluble* games – games where a combination of players' Stackelberg strategies is a Nash equilibrium of the game which, for every player, is associated with a strictly higher personal payoff than any other Nash equilibrium of the game (Colman and Bacharach 1997, Colman and Stirk 1998). The theory offers no definitive theoretical predictions concerning outcomes of games with multiple payoff dominant Nash equilibria. In addition, it does not offer an explanation of how players coordinate actions in games where players have conflicting preferences over the Nash equilibria of the game. For example, in the coordination game depicted in Figure 1.5, the Stackelberg-reasoning decision-makers should end up playing strategy profile (u, r) : The Stackelberg-reasoning row player should choose *u* due to expectation that the column player will respond by playing *l*, while the Stackelberg-reasoning column player should choose *r* due to expectation that the row player will respond by playing *d*. According to the proponents of the theory, the players should not be using Stackelberg heuristic in such cases, which means that the theory cannot be criticized for suggesting that players are using a form of strategic reasoning which guarantees a coordination failure. Yet it can be criticized on the grounds that other descriptive theories, such as, for example, team reasoning theory or social convention theory, can explain coordination both in Stackelberg-soluble and in at least

some of the Stackelberg non-soluble games, thus rendering the Stackelberg reasoning model explanatory irrelevant.

If interpreted as an extension of the orthodox reasoning model, the Stackelberg heuristic is conceptually problematic. A player who uses Stackelberg heuristic engages in ‘magical reasoning’ by fallaciously attributing causal power to his or her own choices. In other words, a Stackelberg reasoner takes his or her own strategy choice as evidence of the opponent’s strategy choice, even though a player who understands the structure of the game should know that players’ choices are causally independent (for extensive discussion of this criticism, see Elster 1989 and Bacharach 2006). Stackelberg reasoner’s beliefs about the opponents’ strategy choices are thus inconsistent with his or her beliefs about the structure of the decision problem.

From the perspective of epistemic game theory, Stackelberg reasoning violates the state-action independence principle, since player’s choice is not independent from his or her probabilistic beliefs about the states of the world – the combinations of opponents’ strategy choices³⁶. Thus, Stackelberg heuristic is incompatible with the fundamental principles of orthodox game theory, and so any explanation of social coordination which employs this heuristic cannot be taken to be a natural extension of orthodox game theory³⁷.

1.3.3 Social Conventions Theory

One of the older theories which offers an explanation of how real-world players resolve games with multiple Nash equilibria is the social conventions theory, the different versions of which have been discussed, among others, by Schelling (1960), Lewis (1969), Sugden (1984), Binmore (2005, 2008), Bicchieri (2006) and Gintis (2008). According to Bicchieri (2006), a social convention can be defined as a behavioural rule c followed by convention followers – a subset of the members of the population P who know that a behavioural rule c exists and applies to situations of type G , where G is a coordination game with two or more strict Nash equilibria. The players who conform to behavioural rule c are playing one of the strict Nash equilibria

³⁶For an early formulation of the state-action independence principle in decision theory, see Savage 1954. For an extensive discussion of the role of state-action independence principle in epistemic game theory, see Perea 2012.

³⁷It is important to emphasize that Colman and Bacharach have proposed the Stackelberg heuristic primarily as an empirical hypothesis concerning people’s actual reasoning in games, not as a model of reasoning compatible with the orthodox conception of rationality (see Colman and Bacharach 1997, Colman 1997, Colman and Stirk 1998). Therefore, an argument can be made that a criticism of Stackelberg reasoning for being incompatible with the principles of orthodox game theory is misdirected.

of the game, which means that any unilateral deviation from the convention yields the deviating player a strictly worse payoff. A convention follower has a conditional preference to follow a convention: If a convention follower expects the other players of the game to follow c , then his or her own decision to follow c is an expected payoff maximizing choice. In other words, players' conformity to a convention is modelled as an expected payoff maximizing behaviour (for discussion, see Bicchieri 2006). Conditional preference assumption implies that convention followers' conformity to c relies on a mutual expectation of conformity: Each player expects the opponents to follow the behavioural rule c in situations of type G , and this expectation makes conformity to c an optimal response³⁸.

Arguably the main difference between the various versions of the social convention theory lies with the structure of beliefs which is assumed to be necessary for sustaining players' conformity. In Lewis's (1969) account, a behavioural rule counts as a social convention if (almost) everyone's conformity to behavioural rule c , (almost) everyone's expectation of (almost) everyone's conformity to c and (almost) everyone's conditional preference to conform to c is common knowledge (see Lewis 1969: 60-80). Lewis's account has been criticised for imposing implausibly strict requirements on individuals' beliefs about each other's beliefs and motivations (see, for example, Binmore 2008), although there is a disagreement over whether the criticisms are based on the same conception of common knowledge as Lewis's theory (see Cubitt and Sugden 2003). Subsequent accounts of social conventions attempted to relax the common knowledge assumptions. For example, in Bicchieri's (2006) account, a behavioural rule c is a convention in a population P if there exists a sufficiently large subset $P^f \subseteq P$, such that each individual $i \in P^f$ knows that a rule c exists and applies to situations of type G , and believes that a sufficiently large subset of P are convention followers – individuals who actually follow c in situations of type G (for extensive discussion, see Bicchieri 2006).

To see how the social convention theory can explain coordination in non-cooperative games, consider the following two player Battle of the Sexes

³⁸It is important to note that this definition of convention is not universally accepted. Some theorists have argued that a convention need not be coordination equilibria, and that convention followers' preference for everyone's conformity to a behavioural rule is not necessary for a behavioural rule to be a convention (see, for example, Sugden 1984 and Vanderschraaf 1998). Other authors have argued that conventions should be treated as correlated equilibria of the game – systems of directives which are such that, given the directives of other players, no player has an incentive not to follow his or her directive (see, for example, Skyrms 1996, Vanderschraaf 1995, 2001 and Gintis 2008). In games with multiple Nash equilibria, the set of correlated equilibria is larger than the set of Nash equilibria, which means that the set of possible conventions is even larger than in the models suggested by Lewis (1969) and Bicchieri (2006).

game depicted in Figure 1.6. In this game, each player has to independently

	o	b
o	2, 1	0, 0
b	0, 0	1, 2

Figure 1.6: Battle of the Sexes game

choose either to play b or o . The game has two strict pure strategy Nash equilibria – (b, b) and (o, o) . The third is a mixed strategy Nash equilibrium $(\frac{2}{3}o, \frac{1}{3}b; \frac{1}{3}o, \frac{2}{3}b)$. Both players strictly prefer coordinating their actions on any Nash equilibrium to playing an out-of-equilibrium strategy profile. However, the players also have conflicting preferences over the pure Nash equilibria of the game: The row player prefers outcome (o, o) over outcome (b, b) , while the column player prefers outcome (b, b) over outcome (o, o) . The players are complete strangers who cannot communicate, and the payoff structure of this game offers no cues on how to coordinate their actions.

Suppose that the row player believes that there is a convention to play o in this game. If the row player believes that the probability of the column player being a convention follower is higher than $1/3$, then playing o is a best response. The row player has an incentive to play o , since his or her choice of b yields a strictly worse payoff. If the column player believes that the probability of the row player being a convention follower is higher than $2/3$, then o is column player’s best response as well.

It is easy to notice that social convention theory relies on assumption that convention followers share a belief that the proportion of convention followers in the population is sufficiently high. This is an empirical assumption about individuals’ shared beliefs, which may or may not be true, and which may be difficult to test empirically. One of the attempts to provide a *theoretical* justification of this assumption is based on the principles of evolutionary game theory. In theory, any population should, over time, converge to an evolutionarily stable state in which everyone is playing an evolutionarily stable strategy. A convention can then be interpreted as a system of evolutionarily stable strategies which can resist invasions of ‘mutants’ playing other strategies (for evolutionary accounts of social norms and conventions see, for example, Sugden 1986, Skyrms 1996, Bicchieri 2006 and Alexander 2007). This interpretation offers an explanation of why conventions may persist for long periods of time, and why individuals living in the same population

hold beliefs which sustain their conformity to a convention³⁹.

Since every strict Nash equilibrium of the evolutionary game is an evolutionarily stable state, any game with more than one strict Nash equilibrium has more than one evolutionarily stable state. Depending on the initial population state (the initial proportions of players using different strategies), the population can converge on any evolutionarily stable state. In a game with multiple strict Nash equilibria, each strict Nash equilibrium may become a conventional solution of the game. Since the trajectory of evolutionary dynamics depends on the initial population conditions, none of the multiple possible evolutionarily stable states can be ruled out on the basis of formal analysis alone. Thus, evolutionary game theory does not warrant a conclusion that, for example, playing *hi* in the Hi-Lo game must necessarily be a conventional solution of the Hi-Lo game for every population, since Nash equilibrium (lo, lo) is an evolutionarily stable state of the evolutionary game, and so playing *lo* may also, over evolutionary time, become a convention.

In any game with multiple strict Nash equilibria, players' ability to regularly coordinate their actions on a strict Nash equilibrium can always be explained as individuals' conformity to a some pre-existing convention. For this reason it is difficult to test the social convention theory empirically. Since conventions are usually defined as behavioural rules which govern players' actions in a particular game (i.e. a particular type of social interaction), no general regularities of behaviour across multiple games can be expected, and so social convention theory can only be tested on a case-by-case basis. Assuming that individuals from different populations or even groups *within* the same population may be following different conventions – a possibility the proponents of the social convention theory do not rule out⁴⁰ – virtually any outcome of a gameplay can be explained as a product of the interaction of individuals following (possibly different) conventions. In games with multiple strict Nash equilibria, the question of whether individuals who manage to regularly coordinate their actions on a Nash equilibrium do so by following a convention, or by using some other type of decision-making procedure cannot

³⁹The underlying idea is that if a population reaches an evolutionarily stable state in which everyone plays an evolutionarily stable strategy of a certain population game, then every member of the population who interacts with other individuals is almost guaranteed to observe other individuals playing that evolutionarily stable strategy in a specific type of social interaction, and thus form a belief that there exists a behavioural rule followed by the majority of population in situations of that type (for extensive discussion, see, for example, Bicchieri 2006 and Gintis 2008). This evolutionary justification of the correlated beliefs assumption will be discussed in considerable detail in chapter 4.

⁴⁰See, for example, Bicchieri 2006 and Elster 1989 who discuss a number of examples of norms and conventions followed by specific social groups, but not by the population in general.

be answered solely on the basis of information about the observed behaviour of individuals – information which, by a substantial number of prominent game theorists, is considered to be the only type of evidence which can be legitimately used for testing competing descriptive game theory models (for extensive discussion, see, for example, Binmore 2005, 2009a and Hausman 2012).

However, there are good reasons to believe that experiments which record psychological data above and beyond individuals' choices must be considered in order to resolve this underdetermination problem (Dietrich and List 2016). The amount of available psychological data is limited: So far not enough experiments have been conducted which not only recorded participants' choices, but also attempted to elicit participants' beliefs and motivations. A handful of experiments with coordination and mixed motive games, in which participants were asked to report either the reasons of their choices or their beliefs about opponents' choices, suggest that few people choose strategies as best-response reasoners. Rubinstein and Salant (2016) asked participating decision-makers to report their beliefs about other players' choices either before or after making their own choice. They found a sizeable proportion of decision-makers not to best respond to their elicited beliefs (see Rubinstein and Salant 2016). Some of the available data indicates that a sizeable proportion of individuals justify their strategy choices by appealing to some notion of mutual advantage. These results support the goal-directed reasoning theories which suggest that people resolve non-cooperative games by identifying the mutually advantageous strategy profiles and playing their part in realizing them (for experimental results, see, for example, Colman and Stirk 1998). Currently there seems to be not enough empirical evidence to definitively conclude that people view their own strategy choices in experimental games as conformity with the behavioural rules of society.

Another, yet related, criticism of the social convention theory is that it does not account for the complexity of players' normative attitudes towards what they perceive to be the 'appropriate' solution of the game⁴¹. Social conventions are arbitrary behavioural rules. According to the theory, each convention follower's conformity to a behavioural rule is sustained by his or

⁴¹The term 'normative' in this context should not be associated with ethical considerations, but rather with *prudential* ones. The term 'normative attitude' refers to player's belief that a certain strategy profile *should* or *should not* be the outcome of a game played by *intelligent* players aiming to advance their personal interests. It is important to note, however, that social convention theory has been criticised for its inability to capture the ethical dimension of the non-theoretical *folk* notion of convention. For an extensive discussion, see, for example, Guala 2013.

her expectation that other individuals will follow that rule, not by a belief that a conventional solution of the game is, in some sense, better than other possible solutions. In other words, player's conformity is based on expectation of what strategies the other players will choose, not what strategies they *should* choose. The theory fails to explain why people often have, in Bacharach's (2006) terms, 'high quality intuitions' that some solutions of games are more compelling than others.

For example, the majority of people have a strong intuition that *hi* is the 'obvious' choice in the Hi-Lo game. But the reason of why people recognize *hi* as the obvious choice is, as Bacharach suggests, not their belief that everyone plays *hi*. People expect everyone else to choose *hi* because they believe that an intelligent player who understands the structure of this game should never choose *lo*, since playing *lo* obviously goes against the mutual interest of the players (for extensive discussion, see Bacharach 2006). The presence of such intuitions indicates that people are not only capable of following the social rules, but also of using relatively sophisticated concepts, such as the concept of mutual interest, when reasoning about solutions of non-cooperative games.

1.4 Strategic Reasoning and Mutual Advantage

Several theories suggest that people resolve at least some of the non-cooperative games by adopting a mode of reasoning which allows them to narrow down the set of possible solutions by identifying those solutions which are mutually advantageous to the interacting players. These theories will be the focus of the subsequent discussion.

1.4.1 Coalitional Rationalizability

The standard rationalizability imposes weak restrictions on players' beliefs and, consequently, strategy choices: Common knowledge of rationality allows the players to eliminate non-rationalizable strategies, yet they are left with multiple rationalizable outcomes in every game with multiple Nash equilibria. Coalitional rationalizability (Ambrus 2006, 2009, Luo and Yang 2009) is a more restrictive solution concept for non-cooperative normal form games than standard rationalizability. It rests on assumption that players of non-cooperative games can coordinate their actions by acting as a tacit 'coalition': Each coalition member confines his or her play to a subset of strategies if it is mutually advantageous for coalition members to do so.

This solution concept rests on the generalization of the standard assumption of game theory that every player aims to maximize his or her expected utility by choosing an optimal response to probabilistic beliefs about opponents' strategy choices. In coalitional rationalizability model, each player is also assumed to be looking for implicit 'coalitional agreements' – restrictions on the strategy space. Each 'coalitional agreement' restricts each coalition member's play to a specific subset of strategies, but does not specify play within the set of non-excluded strategies. This means that a coalitional agreement does not always single out a unique strategy profile⁴².

A coalitional agreement is said to be mutually advantageous if, for every coalition member, every best response within agreement (a strategy not excluded by restriction which maximizes players' expected payoff) to some conjecture compatible with the agreement (i.e. compatible with the assumption that opponents do not play strategies excluded by agreement) yields a higher expected payoff than any best response outside the agreement (i.e. a strategy which maximizes players' expected payoff in light of a belief that opponents may choose strategies outside the agreement)⁴³.

Coalitional rationalizability theory suggests that players use a specific reasoning algorithm – an iterative addition of mutually advantageous, or supported, restrictions⁴⁴. In the first step, the players look for supported restrictions given the set of all the strategies available to players. In the next step, the players consider a set of strategy profiles consistent with the supported restrictions added in the first step, and search for further supported restrictions. The process continues until no further restrictions can be added. The remaining set of strategies is the set of coalitionally rational-

⁴²In the original definition of coalitional rationalizability, due to Ambrus (2006), players are assumed to be using only pure strategies. However, Ambrus also shows that coalitional rationalizability can be extended to mixed strategy space. For details, see Ambrus 2006, 2009.

⁴³In other words, if we associate each conjecture with the expected payoff that a best response strategy to that conjectures yields, then players in the group prefer every conjecture compatible with the agreement to any conjecture to which a strategy outside the agreement is a best response.

⁴⁴The original coalitional rationalizability concept rests on assumption that each player is certain that every other player understands the implicit agreements implied by supported restrictions and always complies with them, and that this certainty is common knowledge among the interacting players. In the cautious rationalizability model, it is allowed that each player thinks that with at most probability ϵ other coalition members are playing outside the supported restriction. Ambrus (2009) shows how a set of ϵ -coalitionally rationalizable strategies can be defined. Luo and Yang 2009 offer an epistemic model of coalitional rationalizability, in which players are assumed to be using a Bayes rule in forming expectations concerning opponents' strategy choices (see Luo and Yang 2009).

izable strategies of the game (for a technical discussion, see Ambrus 2006).

In some games with multiple Nash equilibria, coalitional rationalizability singles out a unique solution. For example, consider the Hi-Lo game (Figure 1.4). It is easy to check that $\{lo\} \times \{lo\}$ is a supported restriction. For each player, strategy lo is a unique best response to any conjecture which assigns a probability higher than $2/3$ to the event of the opponent playing strategy lo . Notice that strategy lo yields an expected payoff of at most 1 for each player (when chosen as a best response to a belief that the opponent plays lo with probability 1). Given restriction $\{lo\} \times \{lo\}$, each player assigns probability 1 to the event of the opponent playing strategy hi , and so strategy hi yields each player an expected payoff of 2. Nash equilibrium (hi, hi) is the unique coalitionally rationalizable solution of this game.

As all of the theories mentioned before, coalitional rationalizability cannot explain how people resolve certain games with multiple Nash equilibria, even those which seem to have intuitively obvious solutions and create no problems for real-world decision-makers. For example, consider the following two player ‘Weak’ Common Interest game depicted in Figure 1.7:

	$t1$	$t2$
$s1$	11, 10	0, 0
$s2$	0, 0	10, 10

Figure 1.7: ‘Weak’ Common Interest game

This game has two pure strategy Nash equilibria $(s1, t1)$ and $(s2, t2)$. The third is a mixed strategy Nash equilibrium $(\frac{1}{2}s1, \frac{1}{2}s2; \frac{10}{21}t1, \frac{11}{21}t2)$. At least intuitively, the Nash equilibrium $(s1, t1)$ seems to be an ‘obvious’ solution of this game, since it maximizes both players’ personal payoffs. However, the Nash equilibrium $(s1, t1)$ is not the unique coalitionally rationalizable solution of this game.

To see why this is so, suppose that players consider restriction $\{s2\} \times \{t2\}$. For the row player, strategy $s2$ is a best response to any conjecture which assigns a probability higher than $11/21$ to column player’s choice of strategy $t2$. Strategy $s2$ yields an expected payoff of at most 10 (when chosen as a best response to a belief that the column player chooses $t2$ with probability 1). Given restriction $\{s2\} \times \{t2\}$, the row player assigns probability 1 to the event of the column player choosing $s1$, and so strategy $s1$ yields an expected payoff of 11.

The same analysis yields different results for the column player. Strategy $t2$ is column player’s best response to any conjecture which assigns a

probability higher than $1/2$ to the row player choosing s_2 . Strategy t_2 yields an expected payoff of at most 10. However, given restriction $\{s_2\} \times \{t_2\}$, strategy t_1 also yields an expected payoff of 10, which means that strategy t_1 does not yield a higher expected payoff than every possible conjecture to which strategy t_2 is column player's best response. Therefore, restriction $\{s_2\} \times \{t_2\}$ is not supported, which means that Nash equilibrium (s_2, t_2) is a coalitionally rationalizable outcome. At least intuitively, this result seems problematic.

The iterative addition of supported restrictions becomes a relatively complex reasoning procedure when applied to games with more than two strategies: The identification of supported restrictions requires players to identify best responses to various conjectures concerning opponents' strategy choices and compare expected payoffs. If coalitional rationalizability were to be interpreted as an approximate description of the process of reasoning by which people arrive at their strategy choices, an argument could be made that, due to complexity, real-world players could only apply such a complex reasoning procedure to simple interdependent decision problems. For example, in two strategy games with two players, any restriction considered by the players automatically leaves them with just one internally consistent conjecture concerning opponents' actions – a conjecture which assigns probability 1 to non-excluded strategy. In such cases, the identification of the supported restriction involves a fairly basic expected utility calculation, which should not create any problems for relatively unsophisticated decision-makers.

In addition, coalitional rationalizability relies on each player's high confidence in the opponents' ability to identify supported restrictions and comply with them. It could be argued that player's degree of confidence should depend on the perceived complexity of the game: In simple games, the player can be reasonably expected to be highly confident in the opponents' abilities, while in complex games high confidence should not be expected. This, again, raises doubts whether players should be expected to confine their play to a set of coalitionally rationalizable strategies in complex games.

Despite certain limitations, coalitional rationalizability offers a novel belief formation model that goes beyond the standard conception of individual rationality. The main conceptual innovation is the idea that players may act as a coalition by engaging in mutually beneficial joint actions in strategic situations where no communication or binding agreements are possible. In other words, the players are assumed to be using a specific criterion of mutual advantage in order to eliminate certain strategies that could not be eliminated on the basis of standard individualistic considerations. In some games with multiple Nash equilibria, players' ability to identify mutually

advantageous outcomes resolves the game definitively (e.g. the Hi-Lo game). In other games it does not single out a unique solution, but at least narrows down the set of possible solutions, thus increasing the probability of successful coordination. The concept of mutual advantage, although formulated somewhat differently, plays a central role in two goal-directed reasoning theories – the team reasoning theory and the hypothetical bargaining theory – which will be the focus of this thesis.

1.4.2 Goal-Directed Reasoning and Mutual Advantage

In recent decades, game theorists have given more attention to the idea that players' choices in games may be influenced by *frames* – systems of concepts that players use when thinking about the game. Player's adoption of a specific frame may depend on the structural and contextual features of the game, some of which may not be captured by the mathematical representation of the game, and, according to some theorists, may lie outside of rational evaluation⁴⁵. This realization prompted an emergence of an idea that certain structural and/or contextual features of the interdependent decision problem may prompt player's adoption of a goal-directed reasoning mode. A player who adopts a goal-directed reasoning mode follows a specific set of inference rules in order to reach a conclusion on what action(s) s/he should take in a particular environment in order to achieve a certain goal. More specifically, when the player adopts a goal-directed mode of reasoning in a game, s/he identifies a certain strategy profile(s) as the desirable solution(s) of the game, and uses a set of inference rules to figure out what strategy s/he ought to choose in order to make the attainment of that outcome possible, given his or her beliefs about the structure of the game and opponents' actions. The identification of the modes of goal-directed reasoning adopted by the players can thus explain their actions in certain types of games (for extensive discussion, see Gold and List 2004, Bacharach 2006, Gold and Sugden 2007b, Smerilli 2014).

Two goal-directed reasoning models – the team reasoning theory and the hypothetical, or virtual, bargaining theory – suggest that players of non-cooperative games may adopt a goal-directed reasoning mode, which enables them to identify outcomes which are mutually beneficial for the interacting decision-makers.

⁴⁵The idea of frames have been discussed, among others, by Gauthier 1975, Bacharach 1993, 2006, Sugden 1995, Bacharach and Stahl 2000, Janssen 2001, Gold and Sugden 2007b and Smerilli 2012.

1.4.3 Team Reasoning and Mutual Advantage

The core idea of team reasoning is that certain structural and/or contextual features of games may trigger a shift in individual's mode of reasoning from individualistic best-response reasoning mode into team reasoning mode. When a person is reasoning as an individualistic decision-maker, s/he focuses on the question 'what it is that *I* should do in a game in order to best promote my personal interests?' The answer to this question involves the identification of a strategy which, given player's beliefs about the other players' strategy choices, maximizes his or her expected payoff. When a person adopts a team reasoning mode, s/he focuses on the question 'what it is that *we* should do in order to best promote *our* interests?' The answer to this question involves the identification of a *strategy profile or profiles* leading to the attainment of the best possible outcome for the group of individuals acting together as a team. When the player identifies the best possible outcome or outcomes for a team, s/he can work out the strategy that s/he has to choose in order to make the attainment of the team optimal outcome possible (Bacharach 2006, Gold and Sugden 2007a).

According to Bacharach (2006), a simple team reasoning framework can be defined as a triple $(\mathcal{T}, \mathcal{O}, \mathcal{U})$, where \mathcal{T} is the set of players in a team reasoning mode, \mathcal{O} is the set of feasible strategy profiles (outcomes) and \mathcal{U} is a team objective function representing a shared ranking of feasible strategy profiles. Team reasoning is as a mode of reasoning where each player $i \in \mathcal{T}$ first works out the best strategy profile $\sigma^* \in \mathcal{O}$ for the team, and then plays his or her component σ_i^* of that profile. More complex models of team reasoning have been developed for scenarios where team reasoning players either know that some of the players are not team reasoning, or are uncertain as to whether the other players of a game are team reasoning or not⁴⁶.

The idea of team reasoning has been developed in a number of different ways, producing various (sometimes incompatible) accounts of the principles underlying team play in non-cooperative games. In general, each account tries to address two key questions: why do players adopt this mode of reasoning in non-cooperative games, and what properties an outcome of the game must have in order to be identified by team reasoners as the best

⁴⁶In *circumspect team reasoning* model, which has been suggested by Bacharach, a team reasoner is uncertain as to whether the other players of the game are team reasoning or not. The probability of each player adopting a team reasoning mode is given by an exogenous probability distribution ω , which is assumed to be common knowledge. Bacharach shows how team reasoning players can use this information to determine team-optimal actions. See Bacharach 2006.

outcome for the team.

Several answers have been offered to the first question. According to a view attributed to Bacharach (2006), the adopted mode of reasoning depends on a decision-maker's psychological frame of mind, which, in turn, may depend on a number of circumstantial factors, but needs not necessarily be driven by conscious deliberation. Bacharach suggested a *strong interdependence hypothesis*, according to which team reasoning is most likely to be adopted by players in games with a *strong interdependence property*. These are games in which a Nash equilibrium in pure strategies is Pareto-dominated by some feasible outcome which may or may not be a Nash equilibrium of the game and which can only be attained by players acting together (for extensive discussion, see Bacharach 2006 and Smerilli 2014).

Sugden (2003) suggests that a decision-maker may choose to endorse a particular mode of reasoning, but this choice may lie outside of rational evaluation. In Sugden's version of the team reasoning theory, an individual only plays his or her part in realizing a mutually beneficial outcome if s/he has a reason to believe that the others will play their part in the attainment of that outcome as well. Hurley (2005a,b) defends the view that player's adoption of team reasoning may be a result of conscious and rational deliberation: Individuals may rationally choose to regard themselves as 'members' of a single collective agency, and consciously commit to acting solely on the interests universalizable to their 'membership'.

Even fewer answers to the second question can be found in the literature on team reasoning. There seems to be a consensus that team reasoners search for solutions of games which are, in some sense, beneficial for the players as a group. Only a handful of suggestions of what properties a team optimal outcome should have can be found in the literature. Bacharach mentions Pareto efficiency as the minimal requirement that a team optimal outcome should satisfy. More specifically, a team optimal outcome should be a Pareto optimal strategy profile which is a Pareto improvement⁴⁷ over the Nash equilibrium of a game (see, for example, Bacharach 1999 and 2006).

Sugden (2011, 2015) proposed the notion of *mutual advantage* as another property that a team optimal outcome should have. The idea is that an outcome selected by a team should be mutually beneficial from every team member's perspective. According to Sugden, an outcome is mutually advantageous if each player's personal payoff associated with that outcome satisfies a particular threshold. The suggested threshold is each player's

⁴⁷An outcome x is said to be a Pareto improvement over another outcome y if, in terms of players' personal payoffs, outcome x makes at least one player better-off without making anyone else worse-off.

personal *maximin* payoff – the highest payoff that a player can guarantee to himself or herself, irrespective of what the other players do (see Sugden 2015).

Despite these suggestions, there seems to be no satisfactory answer to the question of how the notion of mutual advantage could be operationalized and, provided that operationalization is possible, applied in the formal analysis of games. In fact, some of the authors suggest that a specification of team interests that applies to a wide range of games may not be possible, since team-reasoning player’s beliefs about mutually advantageous solutions of games may be based on arbitrary conventions (see, for example, Sugden 2011, 2015).

However, if a general representation of team interests is not possible, then a specification of the structural and/or contextual factors which prompt decision-maker’s adoption of the team mode of reasoning may not be possible as well. Without a clear answer to the question of what outcomes the team-reasoning players aim to attain in non-cooperative games, team-reasoning decision-makers cannot be clearly distinguished from those who do not, and so a rigorous empirical test of the team reasoning theory is virtually impossible. Therefore, there seems to be a serious theoretical reason for believing that a question of whether some generalizable principles underlying team-reasoning decision-maker’s identification of mutually advantageous solutions of games can be identified should be the primary focus of research.

The working assumption of this thesis is that certain generalizable principles underlying team-reasoning decision-maker’s identification of mutually advantageous solutions of games can be identified. In chapter 2, I discuss a possible formal representation of team-reasoning players’ interests, developed in collaboration with Karpus (Karpus and Radzvilas 2016), which, a few differences aside, is broadly in line with the notion of mutual advantage suggested by Sugden (2011, 2015). I also discuss some of the conceptual limitations of the team reasoning theory: Its departure from the standard *ex-post* stability concepts of game theory, such as the Nash equilibrium, and implicit motivation-transformation assumptions, which can be viewed as having a negative impact on the explanatory power of the team reasoning theory.

1.4.4 Hypothetical Bargaining

One of the features of the theory of team reasoning is that it moves away from the orthodox notion of individual rationality. Team-reasoning decision-makers are seeking to maximally advance the interests of the team, but the team optimal outcome does not necessarily maximally advance the personal

interests of every individual who acts as a member of a team.

Like coalitional rationalizability, the theory of hypothetical, or virtual, bargaining is an attempt to incorporate the notion of mutual advantage into game theoretic analysis in a way which would be broadly compatible with the orthodox conception of individual rationality. The core idea of hypothetical bargaining is that players of non-cooperative games choose their strategies on the basis of what strategy profile(s) they would agree to play if they could openly bargain – engage in real negotiations, in which each player can communicate his or her offers to other players and receive their counteroffers. Unlike a team-reasoning player whose aim is to maximally advance the interests of the team, hypothetical bargainer is a *self-oriented decision-maker* – an individual who aims to maximally advance his or her personal interests, and only cares about the interests of other players insofar as their actions can promote or hinder the advancement of his or her own personal interests. Like a best-response reasoner, a hypothetical bargainer is assumed to deviate from the agreement in situations where unilateral deviation is personally beneficial.

A player who reasons as a hypothetical bargainer views the set of mixed and pure strategy *profiles* of the game as a set of *possible* agreements. S/he then identifies a set of *feasible* agreements. Each feasible agreement is a pure or mixed strategy profile, such that no player can exploit the other player(s) by deviating from it. The player then identifies a feasible agreement which, s/he believes, the players would agree on playing in open bargaining, and plays his or her part in realizing the agreement, provided that s/he has a reason to believe that the other players are hypothetical bargainers and will carry out their end of that agreement. An agreement identified by hypothetical bargainers is assumed to be the mutually beneficial and agreeable solution of the game (for details, see Misyak and Chater 2014, Misyak et al. 2014).

Like team reasoning, hypothetical bargaining can be viewed as a goal-directed mode of reasoning: A player accepts certain premises about his or her goal in the decision problem (the player conceptualizes a certain non-cooperative game as a bargaining problem which requires a solution), and then follows a set of well-defined inference rules in order to reach a conclusion on what action s/he should take in order to make the attainment of the goal possible (the player identifies the bargaining solution of the game, and what strategy s/he must to play in order to carry out his or her end of the agreement).

It is important to note that Misyak and Chater (2014) suggest their *virtual bargaining* model as a *cognitive model* – an approximately true descrip-

tion of the mental process by which people arrive at their strategy choices. In other words, the proponents of the theory claim that people actually engage in a mental, or ‘virtual’, *simulation* of the open bargaining process in order to resolve non-cooperative games. However, Misyak and Chater only use their model to predict people’s *behaviour* in certain experimental games. The psychological evidence supporting the mental simulation of open bargaining hypothesis has not yet been provided (see Misyak and Chater 2014, Misyak et al. 2014).

The model of virtual bargaining can also be interpreted as a purely *descriptive* model – either as an approximate accurate description of the practical, or goal-directed, reasoning rules that people use when searching for solutions of non-cooperative games or, purely instrumentally, as an approximately accurate description of people’s choices. A commitment to the empirically ambiguous cognitive interpretation of the model does not seem to be necessary. Since the focus of this study is a descriptive rather than a cognitive interpretation of the model, the term *hypothetical bargaining* rather than *virtual bargaining* will be used throughout this thesis.

Hypothetical bargaining is intuitively appealing. In bargaining games where players’ agreements are not binding, the set of feasible agreements is the set of correlated equilibria. A bargaining solution is a correlated equilibrium which satisfies a number of intuitively desirable properties. According to Myerson (1991), a bargaining solution can be interpreted as an expectation of the outcome of the open bargaining process between players of roughly equal bargaining abilities (see Myerson 1991). Therefore, it seems reasonable to believe that certain properties of bargaining solutions that players deem desirable may play a role in the identification of mutually beneficial solutions of non-cooperative games. Conceptual connections between bargaining and equilibrium selection problems in non-cooperative games have been discussed, among others, by Raiffa (1953), Luce and Raiffa (1957), Aumann (1959), Schelling (1960), Myerson (1991), Moreno and Wooders (1996) and Ambrus (2006, 2009).

One of the fundamental questions pertaining to the theory of hypothetical bargaining is what properties a strategy profile must have in order to be identified by hypothetical bargainers as the hypothetical bargaining solution of the game. Misyak and Chater suggest that ‘existing formal accounts of explicit bargaining, such as Nash’s theory of bargaining, while incomplete, are nonetheless useful as a starting point for the analysis of virtual bargaining’ (Misyak and Chater 2014: 4), and use the Nash bargaining solution as an approximation to what hypothetical bargainers would identify as the mutually advantageous and agreeable solution of a non-cooperative game.

In chapter 3 of this thesis, I will argue that the use of the Nash bargaining solution is conceptually problematic, since it is a bargaining solution function which is not sensitive to the relevant information about the possible alternative allocations of players' individual payoff gains, and therefore does not offer a compelling answer to the question of how hypothetical bargainers identify solutions of games with multiple (weakly) Pareto optimal alternatives associated with different allocations of players' personal payoff gains.

In this chapter, I also propose a benefit-equilibrating (BE) hypothetical bargaining solution concept for non-cooperative games, broadly in line with the principles underlying Conley and Wilkie's (2012) ordinal egalitarian solution for Pareto optimal point selection problems with finite choice sets, which can be applied to cases where interpersonal comparisons of players' payoffs are assumed not to be meaningful. I offer both the ordinal and the cardinal versions of this solution concept, discuss their formal properties, and illustrate their applications in the analysis of non-cooperative games with a number of experimentally relevant examples.

Hypothetical bargaining theory has been introduced primarily as a rational social coordination account – a theory which purports to explain how rational social agents manage to coordinate their actions in various coordination games, in which players cannot communicate and no commonly known social rule of behaviour is available. In chapter 4 I will argue that although hypothetical bargaining theory offers a relatively parsimonious explanation of how people identify the mutually beneficial solutions in a large variety of non-cooperative interdependent decision problems, at best it offers only a partial explanation of how people coordinate their actions in non-cooperative games. I will focus on two epistemic limitations of the theory. I will argue that the theory of hypothetical bargaining, if interpreted as a model of rational decision-making, is vulnerable to the *choice rationalization problem*: The model of hypothetical bargaining does not fully account for the structure of beliefs which sustains hypothetical bargainers' motivation to play their part in the implementation of hypothetical bargaining solutions. I will discuss several responses to this problem and their limitations. I will also argue that hypothetical bargaining, if interpreted as a rational social coordination theory, is vulnerable to the problem of common beliefs: The theory cannot account for the structure of beliefs which makes it a functioning social coordination mechanism. I will also discuss several possible responses to this problems. Finally, I will argue that even a fully developed hypothetical bargaining theory might not be able to provide a single generalizable model of players' final choices due to non-uniqueness of hypothetical bargaining

solutions.

1.5 The Structure of the Thesis

Each chapter of this thesis is a self-contained study of a specific conceptual issue of a particular goal-directed reasoning model. In each chapter, occasional references are made to other chapters in places where conceptual differences and similarities between goal-directed reasoning models are being discussed.

In chapter 2 I will discuss a formal representation of team-reasoning players' interests, developed in collaboration with Karpus (Karpus and Radzvilas 2016), which, a few differences aside, is broadly in line with the notion of mutual advantage suggested by Sugden (2011, 2015). In this chapter I will also discuss several conceptual problems of the team reasoning theory.

In chapter 3 I will suggest the benefit-equilibrating (BE) hypothetical bargaining solution for non-cooperative games, discuss its formal properties, and illustrate its application in the analysis of non-cooperative games with a number of experimentally relevant examples. I will argue that a model of hypothetical bargaining based on the benefit-equilibrating solution concept offers a conceptually credible explanation of how self-oriented decision-makers identify the stable mutually advantageous solutions of non-cooperative games.

In chapter 4 I will discuss the epistemic problems of the hypothetical bargaining theory, and argue that it does not offer a simple answer to the social coordination problem.

With chapter 5 I conclude.

Chapter 2

Team Reasoning and the Measure of Mutual Advantage

2.1 Introduction

One of the conceptual limitations of the best-response reasoning model is its inability to offer theoretical predictions of players' actions in games with multiple rationalizable outcomes, even in those which intuitively seem to have an 'obvious' solution. A canonical example of a game where the standard game theoretic analysis based on best-response reasoning yields conclusions which contradict our intuitions is the Hi-Lo game depicted in Figure 2.1:

	<i>hi</i>	<i>lo</i>
<i>hi</i>	2, 2	0, 0
<i>lo</i>	0, 0	1, 1

Figure 2.1: Hi-Lo game

In this game, two players must simultaneously and independently choose either to play strategy *hi* or to play strategy *lo*. If both players choose strategy *hi*, each gets a payoff of 2. If both choose *lo*, each gets a payoff of 1. If one player chooses *hi* while the other one chooses *lo*, both players get a payoff of 0 – the worst payoff attainable in this game.

Both *hi* and *lo* are rationalizable strategies. The game has two pure strategy Nash equilibria – (*hi*, *hi*) and (*lo*, *lo*). The third one is a mixed strategy Nash equilibrium, in which players randomize between pure strategies *hi* and *lo* with probabilities 1/3 and 2/3 respectively. The mixed Nash equilibrium yields each player an expected payoff of 2/3.

Thereby, best-response reasoning identifies a number of rational solutions of this game and, as a result, it does not resolve the game definitively for the interacting decision-makers¹.

It is true that if one player expected the other player to play *lo* with probability higher than $2/3$, then choosing *lo* would be a best response to his or her belief about the other player's choice. In other words, playing *lo* would be the rational thing for a player to do. It would also be rational for a decision-maker to choose strategy *hi* with probability $1/3$ and strategy *lo* with probability $2/3$ if s/he expected the other decision-maker to do the same. For many people, however, both the Nash equilibrium (lo, lo) and the mixed strategy Nash equilibrium do not appear as convincing rational solutions, while the attainment of the Nash equilibrium (hi, hi) , which is unambiguously the best outcome for both players, appears to be an 'obvious' definitive resolution of this game. The perfect alignment of player's personal interests seems to be the primary reason of why people find an idea that a rational player could expect another rational player to choose *lo* in this game not compelling (Bacharach 2006). Experimental results suggest that over 90% of the time people do opt for strategy *hi* in this game².

This prompted the development of the theory of team reasoning, which suggests that certain contextual and/or structural features of games may trigger a shift in decision-makers' mode of reasoning from individualistic best-responding to reasoning as members of a team – a group of individuals who act together in the attainment of some common goal³. With certain formal definitions of the team goal, the theory can be operationalized to select (hi, hi) as the only solution of the Hi-Lo game for those who reason as members of a team. As such, it offers an intuitively compelling explanation of how people resolve certain games with multiple Nash equilibria.

The idea of team reasoning has been developed in a number of different ways, producing various and sometimes even incompatible accounts of the principles underlying team play in non-cooperative games. In general, the

¹To be accurate, best-response reasoning coupled with a common belief about its application can rationalize certain non-Nash-equilibrium outcomes as well (for extensive discussion, see Bernheim 1984, Pearce 1984, and Perea 2012). For the moment, this discussion will be limited to sets of the Nash equilibria. It is important to note, however, that the possibility of best-response reasoning producing a greater number of rationalizable outcomes only makes the ideas discussed in this chapter more relevant.

²See Bardsley et al. (2010) who, among a number of other games, report experimental results from two versions of the Hi-Lo game where the outcome (hi, hi) yields a payoff of 10 while the outcome (lo, lo) yields a payoff of 9 or 1 to both players.

³For early developments of this theory see Sugden (1993, 2000, 2003) and Bacharach (1999, 2006). For some of the more recent work see Gold and Sugden (2007a,b), Sugden (2011, 2015) and Gold (2012).

theory of team reasoning needs to address two important questions: ‘when do people reason as members of a team?’ and ‘what do people do when they reason as members of a team?’. In other words, it needs to suggest why and under what circumstances people may reason as members of a team and, once they do, what it is that they take team interests to be.

There seems to be a consensus that team reasoners search for solutions of games which are, in some sense, beneficial for individuals as a group. However, only a handful of suggestions of what properties a team optimal outcome must have can be found in the literature. What needs further development is a specification of team interests or team goals that applies, if such generalization is possible, across a wide range of games. In this chapter I will predominantly focus on this question, although some tentative ideas concerning the question of what factors may prompt people to reason as members of a team will be suggested towards the end.

In this chapter I will focus on discussing a function of team interests based on the notion of mutually advantageous play discussed by Sugden (2011, 2015), which have been developed in collaboration with Jurgis Karpus (Karpus and Radzvilas 2016). Our proposed function is compatible with the orthodox conception of payoffs in games, which means that its application does not require payoffs to be interpersonally comparable. Thus, our approach differs from those that use aggregative functions to represent the interests of players who reason as members of a team, such as the maximization of the sum or the average of interacting decision-makers’ personal payoffs.

This chapter is structured as follows. In Section 2 I will discuss the theory of team reasoning in more detail and explain how it differs from the standard payoff transformation approach. In Section 3 I discuss a few properties that Karpus and I believe a potential function of team interest ought to satisfy, and show why aggregative functions may be ill-suited for this purpose. In Section 4 I discuss our proposed measures of individual and mutual advantage in games, and present a function of team interests as the maximization of mutual advantage attained by the interacting decision-makers. I will review the function’s prescriptions in a few examples and discuss some of its properties. In this section I will also revisit the topic of interpersonal comparability of payoffs in order to distinguish interpersonal comparisons of utility from the interpersonal comparisons of advantage based on our working definition of the latter. In Section 5 I will present some of our tentative ideas about why and under what circumstances people may reason as members of a team. In section 6 I will discuss our response to the coordination problem which arises in games when our function produces multiple solutions. In sec-

tion 7 I will consider a couple of conceptual criticisms which can be directed not only against our proposed function of team interests, but also against the theory of team reasoning in general. With section 8 I conclude.

2.2 Team Reasoning

2.2.1 What is Team Reasoning?

When an individual reasons as an individualistic best-response reasoner, s/he focuses on the question ‘what it is that *I* should do in order to best promote *my* interests, given what I know/believe about the motivations and actions of others?’. By answering this question, an individualistic best-response reasoner identifies a *strategy* which is associated with the highest expected personal payoff, given his or her beliefs about the actions of others.

When a person reasons as a member of a team, s/he focuses on the question ‘what it is that *we* should do in order to best promote *our* interests?’. The answer to this question identifies a strategy profile – one strategy for each player in a game – that leads to the attainment of the best possible outcome for the group of individuals acting together as a team⁴. As explained by Gold and Sugden, ‘when an individual reasons as a member of a team, she considers which combination of actions by members of the team would best promote the team’s objective, and then performs her part of that combination’⁵ (Gold and Sugden 2007a: 121).

Bacharach (2006) suggests that a simple team reasoning framework can be defined as a triple $(\mathcal{T}, \mathcal{O}, \mathcal{U})$, where \mathcal{T} is the set of players who reason as members of a team, \mathcal{O} is the set of feasible strategy profiles (outcomes) and \mathcal{U} is a team objective function representing a shared ranking of feasible strategy profiles. Team-reasoning decision-maker $i \in \mathcal{T}$ first works out the best strategy profile $o^* \in \mathcal{O}$ for the team and then plays his or her component

⁴Some versions of the theory of team reasoning consider scenarios in which not all individuals reason as members of a team and where this is recognized by the interacting players. In such cases, the answer to the second question identifies a strategy for every player in a game who does reason as a member of a team. (For an overview see Gold and Sugden, 2007a.) Also, as already noted, the answer to the first question may identify more than one strategy for any one player in a game. This can happen with strategy profiles selected for a team too. Such cases will be discussed later.

⁵In Sugden’s (2003, 2011, 2015) version of the team reasoning theory, player’s commitment to playing his or her part in the attainment of a team optimal outcome may be conditional on the assurance that other players are reasoning as members of a team as well.

o_i^* of that profile⁶.

If in the Hi-Lo game the outcome (hi, hi) is identified as the unique team optimal outcome by players who reason as members of a team, team reasoning can be said to resolve this game definitively for those players who endorse it. This would be so if, for example, the team objective function \mathcal{U} was defined as maximizing the sum/average of the interacting decision-makers' personal payoffs. In that case, a decision-maker who endorsed team mode of reasoning would identify strategy profile (hi, hi) as the uniquely optimal outcome for the team and, consequently, would choose hi in order to play a part in the attainment of this outcome.

A known example of a game where team reasoning can be operationalized to prescribe the attainment of an out-of-equilibrium outcome is the Prisoner's Dilemma game depicted in Figure 2.2:

	c	d
c	2, 2	0, 3
d	3, 0	1, 1

Figure 2.2: Prisoner's Dilemma game

In this game, two players simultaneously and independently choose whether to cooperate (play strategy c) or defect (play d). This game has a unique pure strategy Nash equilibrium (d, d) . Strategy d strictly dominates strategy c : It is always optimal to play d , no matter what the opponent does. In terms of best-response reasoning, strategy c is non-rationalizable, and so a rational player should not be expected to play it. Therefore, strategy profile (c, c) is a non-rationalizable outcome. From the perspective of orthodox game theory, cooperation is not rationalizable in the Prisoner's Dilemma game.

If, however, any of the outcomes involving the play of strategy c is associated with the attainment of team's goal, then strategy c can be selected by team-reasoning decision-makers – individuals aiming to play their part in the attainment of team's goal. If, for example, the team's goal were to maximize the average or the sum of individuals' payoffs, then the team would select outcome (c, c) . As such, with reference to results from numerous

⁶In a *circumspect team reasoning* model, Bacharach (2006) considers a scenario where a team reasoner is uncertain as to whether the other players of the game are team reasoning or not. The probability of each player adopting a team reasoning mode is given by an exogenous probability distribution ω , which is assumed to be common knowledge. Bacharach shows how team reasoning players can use this information to determine team-optimal actions. See Bacharach 2006.

experiments showing that in a one-shot version of the Prisoner's Dilemma game people tend to cooperate about 50% of the time, the theory of team reasoning operationalized this way provides a suggestion of why that is the case⁷.

2.2.2 Team Reasoning and Transformations of Payoffs

A number of game and rational choice theorists have stressed the point that the payoff structure of a game has to capture all the motivations relevant for each player's evaluation of the possible outcomes of a game⁸. Without an accurate information of what game people are playing in terms of their true motivations, it would, in most cases, be impossible to make any conclusions as to whether players' choices are rational and why they make the strategy choices that they do. To see this, consider the monetary Prisoner's Dilemma game depicted in Figure 2.3(a):

	<i>c</i>	<i>d</i>	
<i>c</i>	£2,£2	£0,£3	
<i>d</i>	£3,£0	£1,£1	

(a)

	<i>c</i>	<i>d</i>
<i>c</i>	2, 2	3, 0
<i>d</i>	0, 0	1, 1

(b)

Figure 2.3: Monetary Prisoner's Dilemma game (a) and its transformation (b)

Suppose that the row player is an altruist when it comes to decisions involving money: S/he always strives to maximize the other individual's monetary gains. The column player, on the other hand, is not an altruist, yet s/he is averse to inequitable distributions of monetary gains among the interacting decision-makers. Suppose that his or her inequity aversion is such that any unequal distribution of monetary gains is just as good for him or her as personally gaining nothing. A correct representation of such decision-maker's motivations transforms the monetary Prisoner's Dilemma game into the game shown in Figure 2.3(b). In the transformed game, there

⁷In numerous studies it has been observed that on many occasions cooperation tends to decrease with repetition (that is, when people play the Prisoner's Dilemma game a number of times in a row). This result will be discussed later. See Ledyard 1995 and Chaudhuri 2011 for surveys of experiments with public goods games, which involve more than two players but are otherwise similar in their structure to the two-player Prisoner's Dilemma game.

⁸See, for example, Binmore 1992, 2009a and Hausman 2012.

is one pure strategy Nash equilibrium (c, c) , which means that strategy profile (c, c) is a rationalizable outcome: The players would cooperate if they both were individualistic best-response reasoners and expected each other to play strategy c .

If a theorist were to analyze the observed cooperation using a monetary representation of the game, s/he could reach an incorrect conclusion that players' choices were irrational or that they were reasoning as members of a team.

This raises the question of whether a shift in player's mode of reasoning from individualistic best-response reasoning to reasoning as a member of a team could, in a similar way as in the case of inequity aversion and altruism discussed above, be represented as a transformation of payoffs reflecting decision-makers' motivation to attain the team's goal. To see that this could not be achieved, consider, again, the Hi-Lo game (Figure 2.1). Suppose that team-reasoning decision-makers view outcome (hi, hi) as the best outcome for a team. The outcome (lo, lo) is deemed to be the second-best. The outcomes (hi, lo) and (lo, hi) are deemed the worst. Replacing any of the two players' original payoffs in each cell of the game matrix with those corresponding to the team's ranking of the four outcomes does not change the payoff structure of the original game. As a result, the set of Nash equilibria and, hence, individualistically rational solutions of the transformed game would be the same as the set of rational solutions of the original game. It is exactly the result that the theory of team reasoning was developed to contest. The key difference between the team reasoning theory and the descriptive accounts of transformations of players' personal payoffs based on best-response reasoning is that individualistic best-response reasoning prescribes evaluating and choosing a particular strategy based on the expected personal payoff associated with that strategy, whereas team reasoning prescribes evaluating outcomes of a game from the perspective of a team (given the interacting players' personal preferential rankings of those outcomes) and then choosing a strategy which is an element of the optimal outcome (i.e. strategy profile) for the team. As such, the motivational shift that takes place with a switch from one mode of reasoning to another cannot be captured by transforming players' payoff numbers in the cells of considered game matrices⁹.

Another important question pertaining to the theory of team reasoning is whether a shift in player's mode of reasoning from individualistic best-

⁹The fact that standard transformations of payoffs cannot be used to represent a motivational shift associated with a switch from individualistic reasoning to team reasoning can also be viewed as a more general indication that standard payoff transformation approach cannot be used to capture certain types of relevant motivations of the interacting players.

response reasoning mode to team reasoning mode, or vice versa, changes not only the way in which a decision-maker reasons about which course of action it is best to take, but also the way s/he personally values the outcomes of a game. In other words, team reasoning theory raises an important question of whether a shift in player's mode of reasoning transforms the payoff structure of the game itself. In the formal representation of mutual advantage developed in collaboration with Karpus (2016), we assumed that this does not happen. We believe there to be reasonable grounds for taking this approach. If a shift in decision-maker's mode of reasoning may change the way in which a decision-maker personally values the possible outcomes, then interactions between individuals could become games of incomplete information about the payoff structure of the game. In other words, team-reasoning decision-makers may become uncertain as to what game they are playing. The solutions prescribed by best-response reasoning, team reasoning, and other possible modes of reasoning (e.g. regret minimization and *maximin* reasoning) in games of incomplete information would often be different from those in complete information games. In a game of incomplete information, decision-makers' strategy choices would depend not only on their modes of reasoning, but also on their probabilistic beliefs about the payoff structure of their interaction (for extensive discussion, see chapter 1). The theory of team reasoning was originally developed to resolve complete information games in which orthodox best-response reasoning model was deemed to produce inadequate conclusions. For example, in the case of the Hi-Lo game, the theory was meant to resolve precisely this complete information game with its particular payoff structure¹⁰. Because of the aforementioned considerations, our proposed notion of mutual advantage is based on assumption that players' personal payoffs are common knowledge and a shift in decision-makers' reasoning mode leave their personal payoffs unchanged.

However, later I will argue that team reasoning approach allows for modelling of team-reasoning decision-makers' incentives to play their part in the attainment of the team's goal as being independent from decision-makers' personal incentives that motivate their actions before a shift from individualistic to team mode of reasoning occurs, and that this means that team reasoning may not always avoid making implicit personal motivation transformation assumptions.

The last point to note is that the payoffs attained by the members of a team are assumed not to be transferable. If players were able to share their gains with others, such strategies and the associated payoff distributions

¹⁰It is an important question whether the assumption of common knowledge of payoffs in games is too strong. For extensive discussion, see chapter 1.

would have to be included in representations of their strategic interactions.

2.3 Team Interests and the Notion of Mutual Advantage

2.3.1 Self-Sacrifice and Mutual Advantage

There seems to be a consensus that team reasoners search for solutions of games which are, in some sense, beneficial for the players as a group. However, only a handful of suggestions of what properties an outcome must have in order to be identified by team-reasoning decision-makers as the best outcome for a team can be found in the literature. Bacharach (1999, 2006) mentions Pareto efficiency as the minimal requirement that a team optimal outcome should satisfy. More specifically, a team optimal outcome should be a Pareto optimal strategy profile which is a Pareto improvement¹¹ over the Nash equilibrium of a game. A function which satisfies this criterion is the maximization of the average or the sum of the interacting decision-makers' payoffs. As an example, it has been discussed in some of the earlier developments of the theory (see, for example, Bacharach, 1999), as well as in some of the more recent papers (see, for example, Colman et al., 2008, 2014, and Smerilli, 2012). Although it is used merely as an illustration of how the notion of team-optimal outcome could be incorporated into formal analysis of games, the function is able to offer intuitively compelling theoretical predictions in a number of experimentally relevant games. It is easy to see that in the Hi-Lo (Figure 2.1) and the Prisoner's Dilemma (Figure 2.2) games it selects the outcomes (hi, hi) and (c, c) respectively.

If team reasoning is to be interpreted as a mode of reasoning which does not change the original motivations of the interacting individuals, one of the intuitively undesirable features of the aforementioned function is that in some games it singles out outcomes which require one or more players to completely sacrifice their personal interests for the benefit of others. Consider a slight variation of the Prisoner's Dilemma game depicted in Figure 2.4. In this version of the Prisoner's Dilemma game, the aforementioned function selects strategy profile (d, c) as the unique team-optimal outcome. As such, it would require the column player to completely sacrifice his or her personal interests for the benefit of the row player alone. This example

¹¹An outcome x is said to be a Pareto improvement over another outcome y if, in terms of players' personal payoffs, outcome x makes at least one player better-off without making anyone else worse-off.

	<i>c</i>	<i>d</i>
<i>c</i>	2, 2	0, 3
<i>d</i>	5, 0	1, 1

Figure 2.4: A variation of the Prisoner’s Dilemma game

suggests that Pareto optimality is not a sufficient property to define team optimal outcomes.

Sugden (2011, 2015) suggested the notion of *mutual advantage* as another property that a team optimal outcome should have. The idea is that an outcome selected by a team should be mutually beneficial from every team member’s perspective. Although he does not present an explicit function, Sugden proposes to define an outcome as mutually advantageous if each decision-maker’s personal payoff associated with that outcome satisfies a particular threshold. The suggested threshold is each player’s personal *maximin* payoff – the highest payoff that a player can guarantee to himself or herself, irrespective of what the other players do. For example, in the Hi-Lo game (Figure 2.1) the *maximin* payoff for both players is 0. In the Prisoner’s Dilemma games depicted in Figures 2.2 and 2.4 the *maximin* payoff for both players is 1.

For Sugden, an outcome of a game is mutually beneficial if everyone’s *maximin* threshold is met and each player participates in the attainment of that outcome (that is, in the attainment of personal payoffs associated with that outcome). According to this definition, however, all the Nash equilibria in the Hi-Lo game are mutually beneficial. This is because every equilibrium yields each player a higher payoff than his or her personal *maximin* payoff, and each player’s strategy associated with a particular equilibrium is necessary for the attainment of those payoffs. Thus, the above definition of mutual advantage does not exclude Pareto inefficient outcomes and, by itself, it does not suggest of how further ranking of mutually advantageous outcomes could be established.

2.3.2 Interpersonal Comparisons of Payoffs

Another problematic property of any function which aggregates decision-makers’ personal payoffs is that it requires them to be interpersonally comparable. Such comparisons go beyond the standard assumptions of expected utility theory, which make numerical representations of individuals’ preferences possible, but do not automatically grant their interpersonal compara-

bility. For example, the von Neumann and Morgenstern utility representation of decision-makers' preferences in the Prisoner's Dilemma game depicted in Figure 2.4 allows us to say that the row player prefers the outcome (c, c) over the outcome (d, d) , or that s/he prefers the outcome (d, c) over the outcome (c, c) more than s/he prefers the outcome (d, d) over the outcome (c, d) . It does not, however, allow us to say that the row player prefers the outcome (d, c) over the outcome (c, c) more than the column player prefers the outcome (c, d) over the outcome (d, d) .

This is a consequence of the properties of the von Neumann and Morgenstern utility representation of decision-makers' preferences, which can be derived from the axioms of the expected utility theory, but is unique only up to positive affine transformations: If u is a function representing decision maker's preferences over choice options, then so is any function $u' = au + c$, where $a > 0$ and c are constants (for a detailed discussion of why this is so see, for example, Luce and Raiffa 1957). One of the implications of this result is that the payoff structure of the Prisoner's Dilemma game depicted in Figure 2.2 represents exactly the same preferences of decision-makers as, for example, the Prisoner's Dilemma game depicted in Figure 2.5¹²:

	c	d
c	6, 4	0, 5
d	9, 2	3, 3

Figure 2.5: Another representation of preferences in the Prisoner's Dilemma game.

Because of this, a formal representation of team interests should be expected to select the same outcome under both representations of decision-makers' preferences (that is, be invariant under positive affine transformations of players' payoffs). A function which maximizes the sum or average of the interacting players' payoffs, however, selects different outcomes: outcome (c, c) in a game depicted in Figure 2.2 and outcome (d, c) in a game depicted in Figure 2.5.

Any team reasoning model based on a function which aggregates decision-makers' personal payoffs will be applicable to cases where interpersonal com-

¹²The row and the column players' payoffs in the Prisoner's Dilemma game depicted in Figure 2.5 are derived using the following positive affine transformations of their payoffs in a game depicted in Figure 2.2:

$$u'_r = 3u_{row}$$

$$u'_{col} = u_{col} + 2$$

parisons of payoffs are meaningful. However, an interpersonal comparability of payoffs assumption goes beyond the principles of expected utility theory and thus requires a separate justification¹³. In the absence of a compelling theory justifying such comparisons, most of decision and game theorists assume that payoff numbers are not interpersonally comparable.

In the following section I will present a formal characterization of mutual advantage developed in collaboration with Karpus, which, we contend, is applicable in both cases. I will first present our suggested function for cases where decision-makers' payoffs are assumed not to be interpersonally comparable. I will later explain how, we suggest, the function could be modified to be applicable to cases where interpersonal comparisons of payoffs are possible.

2.4 Team Interests as the Maximization of Mutual Advantage

Let us return to Sugden's (2011, 2015) notion of mutually beneficial outcomes. By itself, the notion of a mutually beneficial outcome does not suggest how much of mutual benefit is gained. In order to make comparisons of outcomes in terms of mutual advantage, we need measures of individual and mutual advantage. In a joint work with Karpus, we propose the following definitions:

Individual advantage: An outcome of a game is individually advantageous to a particular player if that player's attained personal payoff is higher than his or her reference point — a payoff level from which the advantage to that player is measured. The *level of individual advantage* gained is the extent by which that outcome advances the player's personal payoff from his or her reference point relative to the largest advancement possible, where the latter is associated with the attainment of an outcome that s/he prefers the most in the game.

Mutual advantage: An outcome of a game is mutually advantageous to the interacting players if each player's attained personal payoff is higher than his or her reference point — a payoff level from which the advantage to that player is measured. The *level of mutual advantage* associated with an outcome is the largest level of individual advantage that is gained by every player.

¹³This point is also discussed in Sugden 2000.

By definition, the maximum level of individual or mutual advantage is 1. To avoid working with decimals when discussing examples, we can express the levels of individual and mutual advantage in percentage terms, which simply means that the values are multiplied by a factor of 100. For example, if, in a two-player game, both players' reference points are associated with a payoff of 0 and their most preferred outcomes with a payoff of 100, a particular outcome associated with payoffs of 30 and 20 to the two players is said to yield 20 units of mutual advantage. The additional 10 units of individual advantage to one player is not mutual. In other words, a level of individual advantage is simply a percentage of the maximum level of individual advantage attainable to a player in a game (which, by definition, is 1), relative to his or her reference point. A level of mutual advantage is the largest percentage of the maximum attainable individual advantage that is gained by every interacting decision-maker.

Note that our proposed definition of individual advantage is simply a decision-maker's personal payoff when his or her payoff function is normalized so that the most personally preferred outcome of a game is assigned the payoff value of 1, while his or her reference payoff is assigned a payoff value of 0. This can always be achieved by applying an appropriate positive affine transformation of that player's original payoffs¹⁴.

Given the measures of individual and mutual advantage, team interests can now be defined as the attainment of outcomes associated with the maximum mutual advantage. In line with Sugden's (2015) suggestion, we imposed a constraint that each player's personal payoff should be at least as high as his or her personal *maximin* threshold. As for Sugden (2015), this restriction is motivated by the assumption that any team play driven by a joint pursuit of team interests should at least yield each member of a team a payoff that s/he can guarantee to himself or herself individually. Notice that since the level of mutual advantage is the largest level of individual advantage that is gained by every player, this is identical to the *maximization of the minimum level of individual advantage across the interacting players*¹⁵.

¹⁴For an extensive discussion of the 0-1 normalization, or Raiffa normalization, see section 2.4.7.

¹⁵There is a connection between our suggested definition of mutual advantage and Gauthier's (2013) ideas on rational cooperation. For Gauthier, rational cooperation in games is associated with the attainment of Pareto efficient outcomes. His proposal, similarly as here, is to maximize the minimum level of personal gains across players relative to some thresholds below which the players do not cooperate. Gauthier does not specify further what these thresholds are, and he hints at justifying rational cooperation based on the idea of 'social morality'. Somewhat differently from Gauthier's suggestion, cooperative play in our model is explained as a result of interacting players' attempts to resolve games in a mutually advantageous way.

2.4.1 Formalization

For a formal presentation of the proposed function of team interests, let $I = \{1, \dots, m\}$ be a finite set of m players and S_i be a set of pure strategies available to player $i \in I$. A pure strategy outcome is defined as a strategy profile $\mathbf{s} = (s_1, \dots, s_m)$, where $s_i \in S_i$ is a particular pure strategy of player $i \in I$. Let $\mathbf{S} = \times_{i \in I} S_i$ denote the set of all possible pure strategy profiles in a game, and $u_i : \mathbf{S} \rightarrow \mathbb{R}$ denote a payoff function that maps every pure strategy profile to a personal payoff for player $i \in I$. A mixed strategy of player $i \in I$ is a probability distribution over S_i . Let Σ_i be a set of all such probability distributions and $\sigma_i \in \Sigma_i$ be a particular mixed strategy of player $i \in I$, where $\sigma_i(s_i)$ denotes probability assigned to $s_i \in S_i$. A mixed strategy outcome (henceforth, outcome) is defined as a strategy profile $\sigma = (\sigma_1, \dots, \sigma_m)$. Let $\Sigma = \times_{i \in I} \Sigma_i$ be the set of all possible mixed strategy profiles and $u_i(\sigma) = \sum_{\mathbf{s} \in \mathbf{S}} (\prod_{i \in I} \sigma_i(s_i)) u_i(\mathbf{s})$ be the expected payoff of player $i \in I$ associated with a mixed strategy profile $\sigma \in \Sigma$.

In what follows, the function is presented for a case when any mixed strategy play is possible.¹⁶ (For a version of the function when only pure strategies are considered simply replace σ_i , σ , Σ_i , and Σ with s_i , \mathbf{s} , S_i , and \mathbf{S} .) Let $u_i^{max} := \max_{\sigma \in \Sigma} u_i(\sigma)$ denote player i 's personal payoff associated with his or her most preferred outcome, let u_i^{ref} denote i 's reference payoff from which individual advantage to i is measured, and let $u_i^{maximin} := \max_{\sigma_i \in \Sigma_i} \{\min_{\sigma_{-i} \in \Sigma_{-i}} u_i(\sigma)\}$ denote i 's *maximin* payoff level in the game (where $\sigma_{-i} \in \Sigma_{-i}$ denotes a combination of strategies of all players other than i).

For any game where $u_i^{max} \neq u_i^{ref}$, the level of individual advantage of player $i \in I$ associated with a particular strategy profile $\sigma \in \Sigma$ can be defined as follows:

$$u_i^l(\sigma) = \frac{u_i(\sigma) - u_i^{ref}}{u_i^{max} - u_i^{ref}}. \quad (2.1)$$

(Notice that if i 's payoff function u_i is normalized so that $u_i^{max} = 1$ and $u_i^{ref} = 0$, then $u_i^l(\sigma) = u_i(\sigma)$).

For any game where $u_i^{max} \neq u_i^{ref}$ for every $i \in I$, the level of mutual

¹⁶In this paper we focus on one-shot interactions. There is a division of opinion on whether mixed strategy play makes sense in such cases. Perea, for example, refers to mixed strategies in one-shot games as useful ‘‘theoretical objects’’, but ‘not something that people actually use in practice’ (Perea 2012: 32). The function of team interests which is suggested here can be used with and without mixed strategy play. In some games its prescriptions will differ in the two cases and we will indicate this when discussing examples.

advantage of the set of players $I = \{1, \dots, m\}$ associated with $\sigma \in \Sigma$ can be defined as follows:

$$u^\tau(\sigma) = \min_{i \in I} u_i^l(\sigma). \quad (2.2)$$

The proposed function of team interests $\tau : \mathcal{P}(\Sigma) \rightarrow \mathcal{P}(\Sigma)$, where $\tau(\Sigma) = \Sigma^\tau \subseteq \Sigma$, selects a subset from the set of all the possible strategy profiles in a game, such that each selected strategy profile maximizes the level of mutual advantage to the interacting players, subject to each player's personal payoff being at least as high as his or her *maximin* payoff level in the game. Formally, each $\sigma \in \Sigma^\tau$ is such that

$$\sigma \in \arg \max_{\sigma \in \Sigma} u^\tau(\sigma), \text{ subject to } \forall i \in I : u_i(\sigma) \geq u_i^{\maximin} \quad (2.3)$$

or, inserting equations (2.1) and (2.2) into equation (2.3),

$$\sigma \in \arg \max_{\sigma \in \Sigma} \left\{ \min_{i \in I} \frac{u_i(\sigma) - u_i^{ref}}{u_i^{max} - u_i^{ref}} \right\}, \text{ subject to } \forall i \in I : u_i(\sigma) \geq u_i^{\maximin}. \quad (2.4)$$

If $u_i^{ref} = u_i^{max}$ for some $i \in I$, then u_i^l is undefined and there is no outcome which is individually advantageous to $i \in I$. Consequently, there is no outcome which is mutually advantageous for the set of players $I = \{1, \dots, m\}$ as a group, and so $\tau(\Sigma) = \emptyset$. To summarize,

$$\tau(\Sigma) = \begin{cases} \Sigma^\tau \neq \emptyset & \text{when } u_i^{ref} \neq u_i^{max} \forall i \in I \\ \emptyset & \text{otherwise} \end{cases}. \quad (2.5)$$

2.4.2 Two Properties of the Function

Two properties of τ can be derived without further specification of u_i^{ref} . First, provided there is a finite number of players and pure strategies available to each player (an assumption on which the whole discussion in this chapter is based), the set of strategy profiles selected by τ is nonempty, unless for any player in a game the reference payoff u_i^{ref} is the same as the payoff associated with that player's most preferred outcome, u_i^{max} , in which case it is empty. This is so because, for every player i in a game, there always exists at least one *maximin* strategy $\sigma_i^{\maximin} \in \arg \max_{\sigma_i \in \Sigma_i} \{\min_{\sigma_{-i} \in \Sigma_{-i}} u_i(\sigma)\}$, such that $u_i(\sigma_i^{\maximin}, \sigma_{-i}) \geq u_i^{\maximin}$. As such, there is at least one strategy profile $\sigma^{\maximin} = (\sigma_1^{\maximin}, \dots, \sigma_m^{\maximin})$, such that $u_i(\sigma^{\maximin}) \geq u_i^{\maximin}$ for every player $i \in I$, which satisfies the constraint defined in (2.3) and (2.4). So long as $u_i^{ref} \neq u_i^{max}$ for every $i \in I$, since the function τ selects strategy profiles associated with maximum mutual advantage from those

that satisfy the above constraint, and since there is at least one strategy profile that satisfies it, it follows that Σ^τ is nonempty. If, on the other hand, $u_i^{ref} = u_i^{max}$ for some $i \in I$, then u_i' is undefined for that player. In such cases $\tau(\Sigma) = \emptyset$. In words, if, in a particular game, there is nothing that is individually advantageous to some player relative to his or her reference payoff, u_i^{ref} , then there can be nothing that is mutually advantageous to all the interacting players as a group.

Another property of τ is that every strategy profile that it selects is efficient in the weak sense of Pareto efficiency.¹⁷ To see this, suppose that τ selects the strategy profile $\sigma^x \in \Sigma$ when there exists another strategy profile $\sigma^y \in \Sigma$, such that $u_i(\sigma^y) > u_i(\sigma^x)$ for every player $i \in I$ (in other words, σ^x is not Pareto efficient in the weak sense). As long as $u_i^{ref} \neq u_i^{max}$ for every $i \in I$, it follows that $\min_{i \in I} u_i'(\sigma^y) > \min_{i \in I} u_i'(\sigma^x)$. Hence, $\sigma^x \notin \arg \max_{\sigma \in \Sigma} \{\min_{i \in I} u_i'(\sigma)\}$, and so $\sigma^x \notin \Sigma^\tau$.

2.4.3 Reference points

Our suggested function can be used with any set of reference points, relative to which the levels of individual advantage are measured. We have proposed three possible reference points. One of them will be used in a review of examples in the next section.

In every game, each of the possible outcomes can be attained via decision-makers' joint actions. Therefore, one possibility to define players' reference points is to set each decision-maker's reference point to be the worst payoff that s/he can attain in a particular game. This will be the payoff associated with each player's least personally preferred outcome:

$$u_i^{ref} = \min_{\sigma \in \Sigma} u_i(\sigma). \quad (2.6)$$

Such an approach, however, may be criticized on the basis that outcomes which are *non-rationalizable* should not be considered on an equal footing with *rationalizable* outcomes when it comes to establishing decision-makers' reference points. In other words, an argument can be made that non-rationalizable outcomes should be left out from the set of outcomes considered as possible reference points. This rationalizability requirement would

¹⁷A strategy profile is Pareto efficient in the weak sense if there is no other strategy profile available that is strictly preferred to it by every player.

exclude outcomes defined in terms of strictly dominated strategies¹⁸. By letting $\Sigma^{br} \subseteq \Sigma$ to denote the set of rationalizable strategy profiles, we can use the following definition of reference points:

$$u_i^{ref} = \min_{\sigma \in \Sigma^{br}} u_i(\sigma). \quad (2.7)$$

However, rationalizability is a concept specifically associated with the best-response reasoning. This prompts the question of whether a criterion based on the orthodox notion of individual rationality should be used for establishing the reference points of a reasoning mode that is not based on individualistic best-response considerations.

We suggested that there are two ways to justify the use of the restricted set $\Sigma^{br} \subseteq \Sigma$ instead of Σ . The first is to argue that decision-makers approach the game as best-response reasoners to begin with and after that, having drawn conclusions about outcomes they could reach through the application of best-response reasoning, they evaluate everything that follows relative to those conclusions, including what is mutually advantageous for them.

The second way to justify this restriction is to argue that decision-makers' reference points, relative to which they evaluate their individual advantage and, subsequently, the mutual advantage associated with each outcome are, essentially, individualistic. These are decision-makers' personal thresholds, such that anything that is preferentially inferior to them is not individually advantageous to them. Hence it may be argued that their establishment should be based on self-oriented individualistic reasoning and, as such, rationalizability is a conceptually sound restriction to impose.

The third approach, which is closest to the definition of mutual advantage suggested by Sugden (2015), is to use each players' *maximin* payoff level as his or her reference point in a game:

$$u_i^{ref} = \max_{\sigma_i \in \Sigma_i} \{ \min_{\sigma_{-i} \in \Sigma_{-i}} u_i(\sigma) \}. \quad (2.8)$$

This definition of a reference point ensures that the *maximin* constraint in the function τ is met automatically. However, a criticism that such a definition of a reference point does not rule out non-rationalizable outcomes could still apply, since strategy profiles associated with decision-makers' *maximin* payoffs may be excluded from the set of rationalizable outcomes.

¹⁸In any two-player game, an outcome is rationalizable if and only if it does not disappear during the process of iterated elimination of strictly dominated strategies. If there are more than two players, however, an outcome that disappears during such elimination is never rationalizable in the above sense, but the converse is not necessarily true: An outcome may survive iterated elimination of strictly dominated strategies, yet nevertheless be non-rationalizable. For a discussion of these results, see Bernheim 1984, Pearce 1984 and Fudenberg and Tirole 1991.

It is important to note that outcomes associated with *maximin* thresholds, when, in relative preferential terms, they are close to players' most preferred outcomes of games, may sometimes serve as definitive solutions of those games when best-response reasoning leads to indeterminacies. That is, *maximin* thresholds themselves may be mutually advantageous. An example of such a game will be provided in the following section.

2.4.4 Examples

This section shows what the function τ selects in a few simple examples. Two of these – the Hi-Lo and the Prisoner's Dilemma games – have already been introduced. The other two are a version of the Chicken game and the High Maximin game depicted in Figure 2.6. In this discussion of examples, the levels of individual and mutual advantage are computed using the second of the three possible reference points discussed in the previous section: $u_i^{ref} = \min_{\sigma \in \Sigma^{br}} u_i(\sigma)$. A detailed derivation of results in the Chicken game will be presented first and the proposed function's prescriptions for the remaining games will be summarized afterwards.

	<i>l</i>	<i>r</i>
<i>u</i>	10, 1	0, 0
<i>d</i>	4, 4	1, 10

(a)

	<i>l</i>	<i>r</i>
<i>u</i>	10, 1	0, 0
<i>d</i>	9, 0	9, 10

(b)

Figure 2.6: Chicken (a) and High Maximin (b) games

There are three Nash equilibria in the Chicken game depicted in Figure 2.6(a): (u, l) , (d, r) , and $(\frac{6}{7}u, \frac{1}{7}d; \frac{1}{7}l, \frac{6}{7}r)$. Notice that the third is a mixed strategy equilibrium, in which row player randomizes between u and d with probabilities $6/7$ and $1/7$, while column player randomizes between l and r with probabilities $1/7$ and $6/7$ respectively. This yields both players an expected payoff of $10/7$.

For both players, the least preferred *rationalizable* outcome from the set Σ^{br} is (u, r) , the *maximin* payoff level is 1 (the lowest possible payoff associated with strategies d and l), and the most preferred outcome yields a payoff of 10. Thus, for each player, $u_i^{ref} = 0$, $u_i^{max} = 10$, and $u_i^{maximin} = 1$. The levels of individual and mutual advantage, u_i^t and u^τ , associated with the four pure strategy outcomes of the game are shown in Table 2.1 (sorted by u^τ). These, as noted earlier, are expressed in percentage terms (u_i^t and u^τ

Outcome	u'_r	u'_c	u^τ
(d, l)	40	40	40
(u, l)	100	10	10
(d, r)	10	100	10
(u, r)	0	0	0

Table 2.1: The levels of individual and mutual advantage associated with each outcome.

are multiplied by a factor of 100). When only pure strategies are considered, the maximally mutually advantageous outcome is (d, l) and, since it satisfies the *maximin* constraint for both players, it is the unique outcome selected by function τ . The result is slightly different, albeit also yielding a unique solution, if mixed strategies are considered as well. In the latter case, the maximum level of mutual advantage is associated with the mixed strategy profile $(\frac{3}{14}u, \frac{11}{14}d; \frac{11}{14}l, \frac{3}{14}r)$, which yields both players an expected payoff (with reference to the payoff structure in Figure 2.6a) of approximately 43.2. The corresponding approximate level of mutual advantage is 43.2, which is higher than the level of mutual advantage associated with the outcome (d, l) . As a result, $\mathbf{S}^\tau = \{(d, l)\}$ when only pure strategies are considered, and $\Sigma^\tau = \{(\frac{3}{14}u, \frac{11}{14}d; \frac{11}{14}l, \frac{3}{14}r)\}$ when mixed strategies are considered as well. Either way, \mathbf{S}^τ and Σ^τ are singletons, which means that the function τ resolves this game definitively for those who reason as members of a team. Note that in both cases τ selects a non-Nash-equilibrium outcome. Results for the remaining three games are summarized in Table 2.2:

Pure strategies alone	u_i^{ref}	u_i^{max}	u_r^{mxm}	u_c^{mxm}	\mathbf{S}^τ
The Hi-Lo	0	2	0	0	$\{(hi, hi)\}$
The Chicken	0	10	1	1	$\{(d, l)\}$
The High Maximin	0	10	9	0	$\{(d, r)\}$
The Prisoner's Dilemma	1	3	1	1	$\{(c, c)\}$
Mixed strategies	u_i^{ref}	u_i^{max}	u_r^{mxm}	u_c^{mxm}	Σ^τ
The Hi-Lo	0	2	$\frac{2}{3}$	$\frac{2}{3}$	$\{(hi, hi)\}$
The Chicken	0	10	1	1	$\{(\frac{3}{14}u, \frac{11}{14}d; \frac{11}{14}l, \frac{3}{14}r)\}$
The High Maximin	0	10	9	$\frac{10}{11}$	$\{(d; [p]l, [1-p]r)\}^*$
The Prisoner's Dilemma	1	3	1	1	$\{(c, c)\}$

*In Σ^τ for the High Maximin $0 \leq p \leq \frac{1}{10}$.

Table 2.2: A summary of the results for the remaining games

The top section of the table shows parameter values and the selected outcomes in \mathbf{S}^τ when only pure strategies are considered ($u_i^{maximin}$ is abbreviated for row and column player as u_{row}^{mxm} and u_{col}^{mxm} respectively). The bottom

section shows these values and the selected outcomes in Σ^τ when mixed strategies are considered as well.

In the case of mixed strategies in the Hi-Lo game, the *maximin* strategy for both players is to randomize between *hi* and *lo* as in the mixed strategy Nash equilibrium. Irrespective of whether mixed strategies are considered or not, however, (hi, hi) is the unique outcome selected by τ .

The three Nash equilibria in the High Maximin game are (u, l) , (d, r) , and $(\frac{10}{11}u, \frac{1}{11}d; \frac{9}{10}l, \frac{1}{10}r)$. The mixed strategy equilibrium yields an expected payoff of 9 and 10/11 to row and column player respectively. In the case of mixed strategies, column player can secure an expected payoff of at least 10/11 by randomizing between *l* and *r* with probabilities 10/11 and 1/11. As in the Chicken game, the output of τ depends on whether mixed strategies are considered or not. If they are, any strategy profile in which the row player plays *d* while the column player randomizes between *l* and *r* with probabilities $0 \leq p \leq 1/10$ and $1-p$ respectively is maximally mutually advantageous and is included in the set Σ^τ , which means that Σ^τ is not a singleton. This still resolves the game definitively for the interacting players, since τ prescribes a unique strategy choice to row player and it is up to column player alone to select any outcome from the set Σ^τ using $0 \leq p \leq 1/10$ of his or her choice. If only pure strategies are considered, τ selects (d, r) . Note that mutually advantageous play in both cases yields the row player his or her *maximin* payoff. As such, this is an example of a case where a personal payoff associated with maximally mutually advantageous outcome(s) is also player's *maximin* level. By contrast, if u_i^{ref} were set for both players to their personal *maximin* payoffs, τ would select the outcome (u, l) , with and without mixed strategy play.

Lastly, in the Prisoner's Dilemma game the function τ selects (c, c) . Parameter values in the Table 2.2 are based on the Prisoner's Dilemma game depicted in Figure 2.2, but the result is the same for all versions of this game discussed in this chapter.

2.4.5 No Irrelevant Player

One of the potentially problematic properties of our suggested definition of mutual advantage is that the function τ does not rule out any of the interacting players in determining the maximally mutually advantageous outcomes. This implies that no player can be ruled out from strategic considerations as being irrelevant. One of the consequence of this definition is that the player who is indifferent between all the outcomes of a game, with $u_i^{ref} = u_i^{max}$, renders $\Sigma^\tau = \emptyset$. This result is consistent with our definition of mutual ad-

vantage: If a player does not care about how an interdependent decision problem is going to be resolved, there is nothing that would give him or her an incentive to seek out a mutually advantageous resolution of a game.

In cases where one or more players have no incentive to play a part in the attainment of a mutually advantageous outcome, the remaining players may consider searching for a mutually beneficial solution on the game by taking into account their predictions of what strategies the non-team-reasoning decision-makers are going to choose. In such cases, a mutually advantageous play is obviously not possible for all players of the game as members of one team. Note, however, that due to a possibility that non-team-reasoning decision-makers' actions may determine the set of outcomes attainable for team-reasoning decision-makers, the latter may not be able to rule out the actions of the former from their strategic considerations.

2.4.6 Independence of Irrelevant Strategies

As the Prisoner's Dilemma example shows, an addition of a strictly dominated strategy to a game can result in changes in the set of outcomes that the team-reasoning decision-makers would identify as being maximally mutually advantageous. This means that even strictly dominated strategies cannot be treated as irrelevant.

Notice, however, that this result is avoided if the outputs of function τ are limited to rationalizable outcomes. This modification requires all the outcomes outside the set $\Sigma^{br} \subseteq \Sigma$ to be excluded from those considered as potential outputs of τ , and the parameters u_i^{ref} with u_i^{max} to be assigned payoff values associated with each decision-maker's least and most preferred outcome or outcomes in $\Sigma^{br} \subseteq \Sigma$.

Returning to examples discussed in section 2.4.4, this modification would not change the results in the Hi-Lo game, yet the results in the Chicken and the High Maximin games would be different: The outcomes selected by τ when mixed strategies are allowed would be the same as those selected in cases where only pure strategies are considered. This is due to mixed strategy outcomes being non-rationalizable. Finally, the modified function would pick no mutually advantageous outcomes in the Prisoner's Dilemma game, since the Nash equilibrium (d, d) is the only rationalizable outcome. Notice that the modified function allows us to rule out all the strictly dominated strategies as irrelevant.

2.4.7 Interpersonal Comparisons of Advantage

As has been noted earlier, we suggested our function as being applicable to cases where interpersonal comparisons of decision-makers' payoffs are assumed not to be possible. The suggested definition of mutual advantage, however, relies on a specific interpersonal comparison: It equates one unit of individual advantage gained by one individual with one unit of individual advantage gained by the other individual. This potentially raises the question of whether the function τ really avoids making interpersonal comparisons of payoffs in the sense in which they are not warranted by the expected utility theory (see subsection 2.3.2). To defend our claim that it does, we propose the following argument.

Our proposed function relies on a 'zero-one rule', which in the game theory literature is often referred to as *Raiffa normalization* (see Raiffa 1953, Luce and Raiffa 1957). Hausman (1995) suggests a version of a 'zero-one rule' as the only legitimate procedure for comparisons of decision-makers' preference satisfaction levels. That is, as a procedure which does not go beyond the principles of the expected utility theory:

'If it is possible to determine a "correct" cardinal and bounded index of preferences for individuals that is unique up to a positive linear transformation, then . . . there is, I contend, one right way to make interpersonal utility comparisons. One should simply compare the following ratios:

$$\frac{U_i(x) - \min U_i}{\max U_i - \min U_i} \quad \frac{U_j(y) - \min U_j}{\max U_j - \min U_j}$$

where $\max U_i$ and $\min U_i$ are the upper and lower limits of Ira's utility function (roughly, the utilities of his best and worst alternatives), $\max U_j$ and $\min U_j$ the limits of Jill's utility function. The so-called "zero-one rule" assigns the utility value of "1" to the tops of everybody's utility functions and "0" to the bottoms.' (Hausman 1995: 480)

Hausman's version of a 'zero-one rule' is very similar to the Raiffa normalization procedure used in our proposed function. The major difference is that in our discussion the minimum is, in line with Raiffa's (1953) suggestion, each player's reference payoff and the maximum is each player's payoff associated with his or her most preferred outcome in a particular game. Hausman notes that he is 'not proposing that one assign zeros and ones to the best and worst of the *feasible* options available in a given decision problem' (Hausman: 1995: 482; Hausman's emphasis). He discusses

the ‘zero-one rule’ as representing ‘people’s “full” preference rankings of at least all the options they have conceived of’ (Hausman 1995: 482). In Hausman’s interpretation, people’s comparisons of preference satisfaction levels are made against a backdrop of a much wider context than some particular decision problem in question.

In principle, our proposed function τ could be used in this setting: The parameters u_i^{ref} and u_i^{max} could be set to represent decision-makers’ least and most preferred conceivable prospects in general. In this case, the extent of mutual advantage would indicate how individually advantageous an outcome is in the context of the interacting decision-makers’ lives in general. This interpretation of mutual advantage, however, would be different from our intended interpretation of mutual advantage as a representation of how mutually beneficial an outcome is within the context of a particular interdependent decision problem, relative to each decision-maker’s most and least preferred outcomes in that game, as well as to what the decision-makers can expect to attain via individual actions.

According to Hausman, the ‘zero-one rule’ does not imply that the suggested minimums and maximums are comparable in terms of individuals’ well-being or any kind of comparable welfare:

‘To question whether, for example, Jill’s bottom might be “lower” than Ira’s bottom is implicitly to reject the notion of utility as merely representing how well preferences are satisfied. If Jill’s and Ira’s preferences are not satisfied to any extent at all, then there is no way that Jill’s preferences could be better satisfied than Ira’s or that Ira’s preferences could be better satisfied than Jill’s. Nothing is relevant to the comparison except the extent to which preferences are satisfied, and “extent to which preferences are satisfied” is simply position in a preference ranking.’ (Hausman 1995: 480-481)

Hausman’s argument seems to apply to the context of interdependent decision problems discussed here. Even if the interacting players’ levels of well-being are not interpersonally comparable, they can still compare the levels to which their personal interests are advanced (i.e. satisfied) within the context of a particular interdependent decision problem that they are aiming to resolve¹⁹.

The ‘zero-one rule’ is, however, often criticized for implicitly ascribing a particular ratio for making interpersonal comparisons of payoffs at the

¹⁹Luce and Raiffa (1957) also defend the 0-1 normalization as a procedure for establishing a comparison of decision-makers’ interests for cases ‘where interpersonal comparisons are not initially meaningful’ (Luce and Raiffa 1957: 154).

point of establishing a common 0-1 scale. Some theorists have argued that 0-1 scale may lead to inappropriate results when used in decision problems where the difference between the best and the worst options is trivial for one individual and a matter of great importance for another. Binmore (2009b) gives the following example:

‘If Eve is a jaded sophisticate who sees [the best option] as only marginally less dull than [the worst option], whereas Adam is a bright-eyed youth for whom the difference seems unimaginably great, what sense does it make to adopt a method of utility comparison that treats the two equally?’ (Binmore 2009b: 550)

Hausman (1995) quotes similar arguments from other scholars’ works. Hammond (1992: 216), for example, questions the validity of such comparisons in situations where one person is ‘undemanding’ while another is ‘greedy’. Sen (1970: 98) suggests that a disabled person’s preference satisfaction may be ‘uniformly lower’ than that of someone who is fully abled. Rawls (1971: 323) argues that a person who is generally ‘pleased with less’ may often (unfairly) appear as more satisfied than somebody else. Griffin (1986: 120) summarizes this line of criticism by stating that ‘[i]t is not the case that we all reach the same peaks and valleys’.

According to Hausman (1995), the aforementioned criticisms are based on a misinterpretation of what exactly the ‘zero-one rule’ is applied to compare. That is, the criticisms implicitly associates individuals’ preference satisfaction levels with some interpersonally comparable objective notion of welfare or well-being. As such, they are based on an implicit assumption that interpersonal comparisons of welfare can be made and that such comparisons may not always coincide with the interpersonal comparisons of levels of preference satisfaction which, according to Hausman, are compatible with the interpretation of preferences based on the expected utility theory.

Hausman’s analysis of the aforementioned criticisms also applies to potential criticisms of the proposed measures of individual and mutual advantage. If in the context of a game someone contemplates the possibility that for one player stakes might be significantly higher than they are for another, that someone makes an implicit assumption that players’ payoffs can be interpersonally compared. The proposed measures of individual and mutual advantage, expressed as levels of relative advancement of players’ personal interests towards their most preferred outcomes of a game, are supposed to be used in situations where such comparisons are not meaningful. The reason of why the suggested measures of individual and mutual advantage can be applied to such cases is the fact that the scales for these measures

can be established from the commonly known and objectively identifiable points in games without any need of interpersonal comparisons of players' attainable payoffs. In order to use the proposed measures, decision-makers need to know each other's preferences over the possible outcomes of games and their reference points, but they need not be able to make any further interpersonal comparisons of their attained well-being in some objectively comparable way. In other words, decision-makers who know the reference points could work out the maximally mutually advantageous outcomes from the information about their cardinal payoffs provided in the payoff matrix, and so our proposed measures can be used in cases where decision-makers have no clue as to what kind of personal motivations those payoff numbers actually represent.

For the aforementioned reasons, the interpersonal comparisons of individual advantage implied by the suggested function τ are different from the interpersonal comparisons considered in the aforementioned criticisms. When players consider mutually advantageous play in games, they simply equate units of measures of their individual advantage – the advancement of their personal interests relative to what each of them deems to be the personally best and the personally worst outcome of a particular interdependent decision problem. Such comparisons can be performed without decision-makers being able to compare units of the attained personal well-being. In order to make interpersonal comparisons of individual advantage in games, all the players need to know is how much a particular outcome is individually advantageous to a player relative to his or her reference point and the most preferred outcome of a game.

However, this is not to say that interpersonal comparisons of payoffs are never possible, or that people never make them when interacting with each other. According to Binmore (2005, 2009b), human beings may have evolved the ability to make such interpersonal comparisons. Therefore, it is important to note that this possibility does not negate the applicability of the proposed function. In situations where interpersonal comparisons are meaningful, the only step that needs to be added before function τ is applied is the rescaling of decision-makers' measures of individual advantage u_i^j . That is, the appropriate scaling factors must be used to equate these units on the basis of how their preference-satisfaction interpersonally compares.

2.5 Mutual Advantage and the Problem of Coordination

Certain versions of the team reasoning theory are considered by theorists as potentially providing an explanation of how people coordinate their actions in non-cooperative games (see, for example, Crawford et al. 2008, Bardsley et al. 2010, Faillo et al. 2013, 2016). In many games, however, the proposed function of team interests τ will select more than one outcome. In some interdependent decision problems, such as the Hi Maximin game depicted in Figure 2.6b, this will not cause a difficulty for players to coordinate their actions. In many cases, however, the non-uniqueness of maximally mutually advantageous outcomes will create considerable coordination problems. This prompts a question of whether the coordination success rate on a particular outcome should be reflected in the measure of the level of mutual advantage associated with that outcome. To answer this question, we considered several games in which our proposed function leaves the team-reasoning decision-makers with a coordination problem.

Consider a version of the extended Hi-Lo game depicted in Figure 2.7(a):

	<i>hi1</i>	<i>hi2</i>	<i>lo</i>
<i>hi1</i>	10, 10	0, 0	0, 0
<i>hi2</i>	0, 0	10, 10	0, 0
<i>lo</i>	0, 0	0, 0	9, 9

(a)

	<i>hi1</i>	<i>hi2</i>	<i>lo</i>
<i>hi1</i>	10, 10	0, 8	0, 0
<i>hi2</i>	8, 0	10, 10	0, 0
<i>lo</i>	0, 0	0, 0	9, 9

(b)

	<i>hi1</i>	<i>hi2</i>	<i>lo</i>
<i>hi1</i>	10 * 10	0, 0	0, 0
<i>hi2</i>	0, 0	10, 10	0, 0
<i>lo</i>	0, 0	0, 0	9, 9

(c)

Figure 2.7: Three versions of the extended Hi-Lo game

There are seven Nash equilibria in this game: three in pure strategies and four in mixed strategies. The four mixed strategy Nash equilibria yield each player an expected payoff of at most 5. The proposed function would

select both the outcome $(hi1, hi1)$ and the outcome $(hi2, hi2)$ as the maximally mutually advantageous solutions of this game. As such, it leaves the interacting team-reasoning decision-makers facing a coordination problem. Since in terms of payoffs both outcomes are indistinguishable (at this point, a possibility of using strategy labels and the positions of outcomes in the payoff matrix as coordination aids will be ignored), the team-reasoning decision-makers who were to attempt to coordinate their actions on one of the two outcomes without communicating with each other could expect to succeed with probability $1/2$. Such a coordination attempt would yield each decision-maker an expected payoff of 5.

The Pareto inefficient outcome (lo, lo) , however, is unique: It is the only outcome in this game yielding each player a payoff of 9. For this reason, it may be beneficial for the interacting team-reasoning players to focus on the attainment of the outcome (lo, lo) rather than to attempt to coordinate their actions on one of the two Pareto efficient yet indistinguishable outcomes, since the former approach would guarantee each player an expected payoff of 9 instead of the maximum payoff of 5 which could be expected from the latter approach.

Bardsley et al. (2010) and Faillo et al. (2016) suggest the idea that the perceived success rates of the attainment of one from a number of indistinguishable outcomes, or *coordination success rate*, should be incorporated into the function of team interests itself, and use such a representation of team interests to interpret the data obtained from a number of experiments²⁰.

Our proposed measure of mutual advantage u^τ can be easily modified to represent the perceived coordination success rates. That is, in games with a number of outcomes which are indistinguishable in terms of mutual advantage, and in the absence of other coordination aids, the original measures of mutual advantage associated with indistinguishable outcomes could be divided by the number of indistinguishable outcomes in question. For example, in the extended Hi-Lo game depicted in Figure 2.7(a), the level of mutual advantage associated with outcomes $(hi1, hi1)$ and $(hi2, hi2)$, given the coordination success rate of $1/2$, would be 50. Notice that the level of mutual advantage associated with the outcome (lo, lo) would, due to its uniqueness, remain to be 90. With the inclusion of the perceived coordination success rate, our proposed function τ would now select (lo, lo) as the uniquely optimal outcome for a team.

²⁰It is important to note that all the games discussed by Bardsley et al. (2010) and Faillo et al. (2016) are such that players' failure to coordinate their actions on one of the available pure strategy Nash equilibria yields each player a personal payoff of 0. This applies to the extended version of the Hi-Lo game depicted in Figure 2.7(a), but is not the case in a few of its variations which will be discussed here.

On the other hand, team-reasoning decision-makers' coordination of actions may often be possible due to properties of the payoff structure other than the aforementioned coordination success rate. For example, consider another version of the extended Hi-Lo game depicted in Figure 2.7(b). The set of pure strategy Nash equilibria of this game is the same as the set of the Nash equilibria of the extended Hi-Lo game depicted in Figure 2.7(a). The original version of the proposed function τ (i.e. the one which does not account for the perceived coordination success rate) would, again, select outcomes $(hi1, hi1)$ and $(hi2, hi2)$ as team optimal. Notice, however, that although from the perspective of the proposed definition of mutual advantage the two outcomes are just as indistinguishable as they were in the game depicted in Figure 2.7(a), in this case the personal payoff of 8 to the row and the column player associated with outcomes $(hi2, hi1)$ and $(hi1, hi2)$ off the matrix diagonal can serve as an aid in players' coordination of their actions in the attainment of outcome $(hi2, hi2)$. The coordination aid in this case is still the payoff structure of the game, but it is not a pair of payoffs associated with any of the outcomes on which the team-reasoning decision-makers aim to coordinate their actions.

Even more importantly, team-reasoning decision-makers' coordination of actions may be possible due to completely arbitrary factors which have nothing to do with the payoff structure of the game. For an example, consider a version of the extended Hi-Lo game depicted in Figure 2.7(c). This game is identical to the extended Hi-Lo game depicted in Figure 2.7(a), with the exception that outcome $(hi1, hi1)$ is marked with a star. The perceived salience of this outcome has nothing to do with the payoff structure of the game, yet this star can potentially serve as a coordination aid. In fact, there is a number of coordination aids that team-reasoning decision-makers could consider choosing: (1) the presence of a star among the outcomes which are maximally mutually advantageous in terms of the original measure u^τ , (2) the absence of a star among the outcomes which are maximally mutually advantageous in terms of measure u^τ , (3) the uniqueness of the payoff pair associated with outcome (lo, lo) . Which one of the potential coordination aids will ultimately be chosen will depend on team-reasoning decision-maker's beliefs concerning the likelihood that one of the multiple available coordination aids will be recognized and considered by others.

The goal of the preceding discussion of examples was to demonstrate that team-reasoning decision-makers' ability to coordinate their actions in games with multiple team optimal outcomes may depend on factors that have nothing to do with how mutually advantageous the outcomes are for the interacting individuals. In addition, due to a possibility of a game hav-

ing multiple coordination aids, team-reasoning decision-makers may face a *second-order coordination problem* – one related to the choice of a coordination aid. It seems that decision-makers’ choice of the coordination aid in such cases will depend on their beliefs about which aids are most likely to be recognized and considered by every interacting individual. These beliefs may largely depend on the interacting decision-makers’ social experiences, cultural backgrounds, social norms and conventions of the society, and other factors which are not related to the payoff structure of the game itself²¹.

Because of the aforementioned reasons, we believe that it is conceptually fruitful to keep the question of which outcomes are mutually beneficial for the interacting team-reasoning decision-makers separate from the question of how team-reasoning decision-makers coordinate their actions in non-cooperate games with multiple team optimal outcomes. In fact, we believe that a separation of the two questions may be useful from the research point of view: In the presence of multiple potential coordination aids, our proposed measure of mutual advantage u^τ may help to identify a subset of coordination aids that team-reasoning decision-makers should be expected to use in order to coordinate their actions, while maintaining their commitment to resolve the interdependent decision problem in a mutually advantageous way.

2.6 The Triggers of Team Reasoning

As has been noted in the introduction, the theory of team reasoning needs to address the question of why and under what conditions real-world decision-makers may reason as members of a team.

So far, several answers to this question have been suggested in the literature. According to a view attributed to Bacharach (2006), the adopted mode of reasoning depends on decision-maker’s psychological frame of mind, which, in turn, may depend on a number of circumstantial factors, but need not necessarily be driven by conscious deliberation. Bacharach suggested a *strong interdependence hypothesis*, according to which team reasoning is most likely to be adopted by players in games with a *strong interdependence property*. These are games in which a Nash equilibrium in pure strategies is Pareto-dominated by some feasible outcome which may or may not be a Nash equilibrium of the game and which can only be attained by players acting together (see Bacharach 2006 and Smerilli 2014).

²¹Bacharach (1993), Bacharach and Bernasconi (1997) and Bacharach and Stahl (2000) developed a formal model — the variable frame theory — for incorporating such considerations into formal representations of coordination problems.

Sugden (2003) suggests that decision-maker may choose to endorse a particular mode of reasoning, but that this choice may lie outside of rational evaluation. Hurley (2005a,b) defends the view that player's adoption of team reasoning may be a result of conscious and rational deliberation: Individuals may rationally choose to regard themselves as 'members' of a single collective agency, and consciously commit to acting solely on the interests universalizable to their 'membership'. A detailed review of these developments will not be provided here (for a survey see, for example, Gold and Sugden 2007a,b), but some tentative suggestions in connection to our proposed function τ will be discussed.

The development of the team reasoning theory was primarily motivated by the fact that the standard best-response reasoning model is unable to single out intuitively compelling solutions in certain types of interdependent decision problems with multiple rationalizable outcomes. We are inclined to believe that a decision-maker who first approaches the interdependent decision problem as a best-response reasoner may switch into team mode of reasoning in situations where the best-response reasoning is unable to resolve the interdependent decision problem definitively. Player's subsequent endorsement of the team mode of reasoning – his or her decision to play a part in the attainment of the maximally mutually advantageous outcome – may depend on a number of factors, such as his or her beliefs about the modes of reasoning endorsed by other decision-makers, as well as beliefs about the outcomes s/he could expect to attain by playing as a best-response reasoner. With respect to our proposed function τ , this provides a reason for considering only rationalizable outcomes as reference points for our proposed measures of individual and mutual advantage: If decision-makers first approach the interdependent decision problems as best-response reasoners, they should base their further evaluations of outcomes on the conclusions which they can draw from the (possibly common) application of the best-response reasoning approach. This could explain why players may switch from individualistic best-response reasoning into team reasoning mode in games with multiple rationalizable outcomes, such as the Hi-Lo game, the Chicken game, and the High Maximin game discussed in section 2.4.4, but it does not explain why some people allegedly reason as members of a team in the Prisoner's Dilemma game.

It seems plausible that a shift in decision-maker's mode of reasoning in the Prisoner's Dilemma and similar games may be triggered by efficiency considerations. However, if the perceived inefficiency (say, in terms of weak Pareto efficiency) of the best-response solution may indeed trigger a shift in decision-makers' mode of reasoning, then such decision-makers may end

up facing two competing definitive resolutions of a game: One based on best-response reasoning and another one based on mutual advantage considerations. A game which is a trivial decision problem from the perspective of orthodox game theory (e.g. Prisoner's Dilemma) may turn into a complicated dilemma about which mode of reasoning a decision-maker should endorse when choosing his or her strategy, with efficient mutually advantageous solution from which the other players may be tempted to deviate pitted against inefficient yet less risky best-response solution. This story seems to fit with the aforementioned findings from experiments with the Prisoner's Dilemma game, which show that cooperation rate is around 50% in one-shot versions of this game²².

Another possibility is that decision-makers first approach the interdependent decision problems as team reasoners. In this case, a team-reasoning decision-maker may be motivated to switch to best-response reasoning mode in situations where a unilateral deviation from the attainment of a mutually advantageous outcome is personally beneficial to him or her, or when team reasoning mode is unable to resolve the game definitively. This can explain why team-reasoning decision-makers may switch to best-response reasoning in certain games, such as the Prisoner's Dilemma, the Chicken game and the High Maximin game (depending, of course, on which outcome of the High Maximin game the team-reasoning decision-makers consider to be mutually advantageous). With respect to our proposed function τ , if decision-makers approach the interdependent decision problems as team reasoners, then their reference points need not be rationalizable outcomes.

However, if in certain games, such as, for example, the Prisoner's Dilemma game, a decision-maker may end up vacillating between the two modes of reasoning when deciding which one to endorse for choosing his or her actions, s/he should engage in a comparison of the perceived advantages associated with both reasoning modes: The advantages associated with team reasoning should be compared with the advantages associated with best-response reasoning²³. It seems reasonable to expect a decision-maker to measure the advantages of team reasoning relative to outcomes which could result from

²²In the Prisoner's Dilemma, a decision to play a part in the attainment of a mutually advantageous outcome is risky compared to a choice of a best-response, since the former decision is only beneficial to a particular player if the other player does likewise. This may explain why, in a repeated setting, the cooperation rate in this game drops in the presence of a few defectors. A decision-maker who first endorses team reasoning, but recognizes that best-response, albeit less efficient when endorsed by everyone, is less risky when it comes to the worst case scenario, may quickly switch to endorsing best-response reasoning after encountering a player who defects.

²³Smerilli (2012) provides a formal model representing the vacillation process between competing modes of reasoning.

the application of best-response reasoning. This provides further grounds for establishing reference points on the basis of rationalizable outcomes of a game.

2.7 Conceptual Limitations

In sections 2.2 to 2.6, I presented a function of team interests developed in Karpus and Radzvilas (2016). In this section I will consider a couple of conceptual criticisms which can be directed not only against our proposed function of team interests, but also against the theory of team reasoning in general.

2.7.1 Implicit Motivation Transformations

In our collaborative work which led to the development of function τ , we primarily focused on the question of what properties an outcome of a game must have in order to be *identified* as mutually advantageous by the interacting team-reasoning decision-makers. As has been pointed out in section 2.2.2, our proposed function is based on assumption that a shift in decision-maker's mode of reasoning from individualistic best-response reasoning to reasoning as a member of a team does not change the way in which that decision-maker personally values the outcomes of the game. Our proposed function operates on the basis of players' payoffs as their are represented in the original game. That is, team-reasoning decision-makers are assumed to identify the team optimal outcomes from the commonly known information about their personal evaluations of the possible outcomes of a game.

As a model which purports to explain how team-reasoning decision-makers *identify* the mutually advantageous outcomes of non-cooperative games, our proposed function and its underlying assumptions are, I contend, rather unproblematic. However, the original aim of the team reasoning theory was to explain people's *actual choices* in certain types of social interactions, including those types of interactions where a mutually advantageous outcome is neither a Nash equilibrium of a game, nor even a rationalizable outcome (e.g. the Prisoner's Dilemma game).

If a strategy profile of a game is not a Nash equilibrium, then, by definition, at least one player can maximize the advancement of his or her personal interests by playing a different strategy than his or her strategy in that profile. In other words, at least one player who expects the other players to play their part in the attainment of an out-of-equilibrium outcome has an opportunity to maximally advance his or her personal interests by playing

a different strategy than the one leading to decision-makers' attainment of that out-of-equilibrium outcome. From the perspective of orthodox game theory, a player who foregoes such an opportunity to advance his or her personal interests can be said to be irrational, provided, of course, that the payoff structure of a game used to analyze decision-maker's actions correctly represents his or her personal motivations (see Binmore 2005, 2009a).

The theory of team reasoning predicts an out-of-equilibrium play in certain types of games. That is, it predicts that at least one team-reasoning decision-maker will play his or her part in the attainment of a mutually advantageous outcome while having an opportunity to choose a strategy leading to the attainment of an outcome which the decision-maker *personally* values more than the team optimal outcome. In other words, team-reasoning decision-makers' attainment of an out-of-equilibrium team optimal outcome implies that at least one decision-maker has foregone an opportunity to attain the outcome which s/he *personally* values more than the team optimal outcome.

From the perspective of orthodox game theory, according to which payoffs effectively represent decision-makers' *choices* over outcomes (that is, the 'all-things-considered' evaluations of outcomes)²⁴, an argument can be made that a team-reasoning decision-maker's choice to play a part in the attainment of a team-optimal outcome while having an opportunity to attain a more personally advantageous outcome is an indication that team-reasoning decision-maker's personal motivations are different from his or her individualistic motivations which are represented by the payoff structure of the original game. More specifically, a decision-maker's team-reasoning-based personal evaluation of outcomes which leads him or her to playing a part in the attainment of a team optimal outcome may be different from his or her best-response-based personal evaluation of outcomes which is represented by the payoff structure of the original game. This would imply that a shift in decision-maker's mode of reasoning from individualistic best-response reasoning to reasoning as a member of a team, and possibly *vice versa*, may transform decision-makers' personal evaluation of outcomes, even if such a transformation cannot be represented with a simple transformation of decision-maker's personal payoffs, similar to the ones which have been discussed in section 2.2.2. This would mean that team-reasoning decision-makers may be playing a game, the structure of which is different from the original game.

This criticism may be viewed as revealing certain explanatory limitations of the team reasoning theory. If it is the case that decision-makers who

²⁴For extensive discussion, see chapter 1.

reason as members of a team may be playing a game with a different payoff structure than the game that those decision-makers would be playing as best-response reasoners, then team reasoning theory can be criticized for not being able to show that team-reasoning decision-makers actually play a part in the attainment of an out-of-equilibrium team optimal outcome. For this reason, it could be criticized for not being able to explain why people cooperate in situations where their personal motivations are identical to, for example, the motivations of players in the Prisoner's Dilemma game.

A criticism that team reasoning involves implicit transformations of decision-makers' personal evaluations of outcomes can also be applied to cases where the team optimal outcome is a Nash equilibrium. Consider the coordination game depicted in Figure 2.8:

	$r1$	$r2$	$r3$
$r1$	10, 8	0, 0	0, 0
$r2$	0, 0	8, 10	0, 0
$r3$	0, 0	0, 0	8, 9

Figure 2.8: Coordination game with a conflict of players' preferences

This game has three weakly Pareto optimal pure strategy Nash equilibria: $(r1, r1)$, $(r2, r2)$ and $(r3, r3)$. Notice that as individualistic reasoners, the players of this game have conflicting preferences over the three outcomes: As an individualistic reasoner, the row player prefers the attainment of outcome $(r1, r1)$ over the attainment of outcomes $(r2, r2)$ and $(r3, r3)$, while the column player prefers the attainment of outcome $(r2, r2)$ over the attainment of outcomes $(r3, r3)$ and $(r1, r1)$ (also notice that the column player prefers the attainment of outcome $(r3, r3)$ over the attainment of outcome $(r1, r1)$).

The aforementioned structure of personal motivations would not be preserved if the interacting decision-makers were to adopt the team reasoning mode and search for a mutually advantageous solution of this game in a way which has been discussed in this chapter. Coordination success considerations aside, our proposed function τ would select all three Nash equilibria as maximally mutually advantageous outcomes. Our proposed function leads to a conclusion that a team-reasoning decision-maker should be indifferent between the three outcomes. That is, an individual whose only motivation is to play a part in the attainment of a maximally mutually advantageous outcome should see all the three outcomes as equally good, since all of them maximize mutual advantage. It can be argued that self-oriented decision-makers whose joint actions were motivated purely by individual advantage maximization

considerations would not endorse such a resolution of this game. In other words, even if self-oriented decision-makers were to search for a mutually advantageous resolution of the aforementioned game, they would, *due to individual advantage allocation considerations*, not be indifferent as to which one of the three weakly Pareto optimal outcomes to implement. The set of outcomes which self-oriented decision-makers would identify as mutually advantageous and *implementable* solutions of a game may thus be different than the set of maximally mutually advantageous outcomes identified by the team-reasoning decision-makers.

This example suggests that team-reasoning decision-makers' actions may be driven by incentives which are different from their individualistic incentives as they are represented by the payoff structure of the original game. It also reveals one of the features of the team reasoning approach in general: It allows for modelling of team-reasoning decision-maker's incentive to play a part in the attainment of the team's goal as being independent from decision-maker's personal incentives that motivate his or her actions before a shift from individualistic best-response reasoning to reasoning as a member of a team occurs.

2.7.2 Stability Issues

The possibility of team-reasoning decision-makers choosing strictly dominated strategies, as well as the possibility that team-reasoning decision-makers may consciously choose strategies leading to out-of-equilibrium outcomes makes it difficult to defend Bacharach's (2006) theory that people's ability to team reason is an evolutionary adaptation. According to Bacharach, people's ability to reason as members of a team in the attainment of efficient outcomes can be explained as an outcome of the group selection process. The idea is that a group of individuals whose members have the ability to coordinate their actions in the attainment of efficient outcomes will have a higher average fitness²⁵ than the group whose members have no such ability (provided, of course, that individuals belonging to a group whose members have the efficient coordination ability interact with each other sufficiently frequently). A group of individuals with efficient coordination ability could thus grow faster than a group of individuals with no such ability. For example, a group of team-reasoning decision-makers who only interact with each other would attain a higher payoff in social dilemmas, such as the Prisoner's

²⁵In standard models of evolutionary game theory, such as those based on replicator dynamics, individual's expected payoff is assumed to represent his or her fitness – the expected number of offspring or imitators who will have individual's trait. For details, see Weibull 1995.

Dilemma game, than a group of best-response reasoners who only interact with each other.

The problem with this evolutionary explanation is that a group of individuals whose members are choosing strategies that do not constitute a Nash equilibrium of the game cannot be evolutionarily stable (for proofs, see van Damme 1987 and Weibull 1995). This means that a group of team reasoners could be invaded by ‘mutants’ – individuals using other reasoning modes, such as, for example, best-response reasoning. Note that best-response reasoners could exploit team-reasoning individuals in many games where team-reasoning individuals’ strategy choices do not constitute a Nash equilibrium, such as the Prisoner’s Dilemma game or the Chicken game. In such interactions, best-response reasoners would get a higher payoff when playing against the team-reasoning individuals than individuals who reason as members of a team. Depending on what games and how frequently a group is playing, best-response reasoners’ ability to exploit team-reasoning individuals may allow them to spread in the group and drive the individuals who reason as members of a team out of the group. From the perspective of evolutionary game theory, an argument can be made that team-reasoning should *at least* not be observed in recurring games where the team optimal outcome is an out-of-equilibrium strategy profile²⁶.

2.8 Conclusion

In this chapter I predominantly focused on discussing a possible formal representation of team-reasoning decision-makers’ interests based on the notion of mutual advantage in games, which has been developed in collaboration with Karpus. We are inclined to believe that the spirit of our working definition of mutual advantage is broadly in line with the notion of mutual advantage suggested by Sugden (2015). Our proposed function of team interests allows us to further discriminate the outcomes which satisfy Sugden’s (2015) criterion of mutual benefit, as well as to define team-reasoning players’ interests as the attainment of an outcome yielding a maximum mutual advantage. I discussed our argument that the proposed function τ is applicable to cases where interpersonal comparisons of the interacting decision-makers’ payoffs are assumed not to be meaningful, as well as to cases where such comparisons are possible.

It is important to emphasize that the proposed model should not be

²⁶For an extensive evolutionary game theoretic analysis of the evolutionary success of cooperative behaviour in social dilemmas, see, for example, Skyrms 1996 or Alexander 2007.

viewed as a complete account of the theory of team reasoning, but merely as an attempt to provide a formal characterization of the properties that an outcome must have in order to be identified as maximally mutually advantageous by players who reason as members of a team. Only a few tentative ideas concerning factors that may motivate decision-makers to adopt the team reasoning mode or to switch between individualistic best-response and team modes of reasoning have been suggested. Further empirical research is required to test the competing theories about the modes of reasoning that people use when interacting with each other. One of the major challenges is the problem of underdetermination: Since decision-makers' actions can often be explained in terms of multiple competing descriptive theories, further empirical research on the team reasoning theory may need to consider a broader evidence base than mere observations of people's behaviour. According to Dietrich and List, our evaluation of theories concerning people's choices in games may need to consider 'novel choice situations, psychological data over and above choice behaviour, verbal reports, related social phenomena, and occasionally (for plausibility checks) even introspection' (Dietrich and List 2016: 273). Elicitation of people's beliefs about each other's actions and further development of such experimental techniques may ultimately be the only viable approach to test the ideas discussed here empirically.

A question of how team-reasoning decision-makers coordinate their actions in games with multiple maximally mutually advantageous solutions also warrants further empirical and conceptual investigation. In fact, since team-reasoning decision-makers' ability to coordinate their actions may depend on arbitrary factors which have nothing to do with payoff structures of games, a single generalizable model of team-reasoning decision-makers' *final* choices may not be possible at all. As has been pointed out by Sugden, mutually beneficial team play may be based on players' conformity 'to complex and sometimes arbitrary conventions that could not be reconstructed by abstract rational analysis' (Sugden 2015: 156).

While a view that decision-makers' final choices in coordination games are often based on arbitrary rules and conventions is almost certainly the right one to hold, I nevertheless believe that some generalizable principles of how the interacting decision-makers identify the mutually advantageous outcomes of games *before* they address the coordination issue can be identified.

However, as I have pointed out in section 2.7, one of the potential problems of the proposed function is that it may involve an implicit result that a shift in decision-makers' mode of reasoning from individualistic best-response reasoning to reasoning as a member of a team involves a transformation of decision-makers' personal incentives: Team-reasoning decision-makers moti-

vated by mutual advantage considerations would not care about the distribution of their individual advantage gains as long as the set of considered outcomes contained only the maximally mutually advantageous outcomes. It could be argued that the proposed function's insensitivity to information about the possible alternative allocations of individual advantage gains makes it unsuitable to represent the notion of mutual advantage which could guide the actions of self-oriented decision-makers whose motivation to engage in joint actions is based purely on individual advantage maximization considerations.

In the next chapter, I will provide an alternative formal characterization of mutual advantage based on the principles of hypothetical bargaining theory. I will argue that hypothetical bargaining theory offers conceptual foundations for an individualistic explanation of how people identify mutually advantageous and implementable solutions of non-cooperative games.

Chapter 3

Hypothetical Bargaining and the Equilibrium Selection In Non-Cooperative Games

3.1 Introduction

A central solution concept of the orthodox game theory is the Nash equilibrium – a pure or mixed strategy profile which is such that no rational player is motivated to unilaterally deviate from it by playing a different strategy. However, at least intuitively certain Nash equilibria are more convincing rational solutions of games than others. Even a very simple game may have a Nash equilibrium which seems unlikely to be played by decision-makers who understand the structure of that game.

Consider a common interest game shown in Figure 3.1, in which two players simultaneously and independently choose between two strategies: The row player chooses between strategy $s1$ and strategy $s2$, and the column player chooses between strategy $t1$ and strategy $t2$. The left and the right number in each cell represents row and column player's payoffs respectively¹.

At least intuitively, the outcome $(s1, t1)$ stands out as an 'obvious solution' of this game: A player who knows the payoff structure of the game should realize that strategy profile $(s1, t1)$ is the best outcome for both players, and that there is no conflict of players' interests in this game. According to Bacharach (2006), such an alignment of players' personal interests in a

¹Unless it is stated otherwise, the payoff numbers in the matrices are the von Neumann and Morgenstern utilities. The payoffs are assumed to represent all the relevant motivations of players, including pro-social preferences, such as inequity aversion, altruism, sensitivity to social norms, and so on.

	<i>t1</i>	<i>t2</i>
<i>s1</i>	4, 4	1, 2
<i>s2</i>	2, 1	3, 3

Figure 3.1: Common interest game

game is the primary reason of why most people have a ‘high-quality intuition’ that strategies *s2* and *t2* are unlikely to be chosen by intelligent decision-makers who understand the structure of this game. Experimental results support this intuition by revealing that over 90% of the time people opt for strategies *s1* and *t1* in this game².

From the perspective of orthodox game theory, there are three rational solutions of this game – two pure strategy Nash equilibria (*s1, t1*) and (*s2, t2*) and a mixed strategy Nash equilibrium ($\frac{1}{2}s1, \frac{1}{2}s2; \frac{1}{2}s1, \frac{1}{2}t2$). Contrary to intuition, the theory does not single out the Nash equilibrium (*s1, t1*) as the unique rational solution of this game.

The reason of this result becomes clear when we look into the model of reasoning which underpins the standard game theoretic analysis. In standard game theoretic analysis of complete information games, players’ rationality and the payoff structure of the game are assumed to be common knowledge³. In orthodox game theory, a player is said to be rational if s/he always chooses a best response – a strategy which, given player’s consistent probabilistic beliefs about the opponents’ strategy choices, maximizes his or her expected payoff. If rationality is common knowledge, then every decision-maker knows that none of the opponents’ will ever choose a non-rationalizable strategy – a strategy which is never a best response for a player, irrespective of what probabilistic beliefs s/he holds about the opponents’ strategy choices. If the payoff structure of the game is also common knowledge among the interacting players, then each player can iteratively eliminate the non-rationalizable strategies, thus identifying the set of *rationalizable* strategies of the game. Each rationalizable strategy is a best response to some *possible* consistent probabilistic belief about the opponents’ strategy choices (for a detailed tech-

²See Colman and Stirk (1998) who, among a number of other games, report results from experiments with the game depicted in Figure 3.1.

³There is an important question of whether the assumption that players’ von Neumann and Morgenstern utilities are common knowledge is not implausibly strong (see chapter 1). This question, however, illuminates an important conceptual problem of game theory in general, not a conceptual problem of the hypothetical bargaining theory in particular. The question of how hypothetical bargaining could be applied to cases where information about players’ cardinal utilities is not available will be addressed in section 3.3.

nical discussion and proofs, see Bernheim 1984 and Pearce 1984).

Every Nash equilibrium is a strategy profile, such that each player's strategy is a best response to a combination of opponents' strategies. By definition, each Nash equilibrium is a *rationalizable strategy profile* – a combination of players' rationalizable strategies. Therefore, every game with multiple Nash equilibria has multiple rationalizable outcomes, which means that at least one player has more than one rationalizable strategy. For example, the Nash equilibrium (s_2, t_2) in the common interest game depicted in Figure 3.1 is, like the Nash equilibrium (s_1, t_1) , a rationalizable outcome: It is rational for the row/column player to play strategy s_2/t_2 if s/he believes that the probability of the opponent playing strategy s_2/t_2 is higher than $1/2$, since in that case strategy s_2/t_2 is the unique best-response. Therefore, all the strategies available to players in the game depicted in Figure 3.1 are rationalizable.

Classical game theory does not offer a model of how rational players form beliefs about each other's rationalizable strategy choices, and therefore cannot answer certain important questions, such as how players coordinate their actions on a Nash equilibrium, or which Nash equilibrium, if any ⁴, will be the most likely outcome of a rational gameplay (Olcina and Urbano 1994).

This prompted the emergence of multiple theories which purport to explain how players resolve games with multiple rationalizable outcomes. One of those theories is the team reasoning theory, which suggests that certain structural and/or contextual features of a game may trigger a shift in player's mode of reasoning from individualistic best-response reasoning to reasoning as a member of a team. A decision-maker who reasons as a member of a team identifies a *strategy profile* that leads to the attainment of the best possible outcome *for the group of individuals acting together as a team*, and works

⁴Aumann and Brandenburger (1995) have established the epistemic conditions of Nash equilibrium, both for two player games and games with $n > 2$ players. In a two player game, if rationality of the players and each player's conjecture (that is, a subjective probability distribution over opponent's strategies) are mutually known, then players will end up playing one of the Nash equilibria of the game. In a game with more than two players, the epistemic conditions of Nash equilibrium are more complicated: Players must have a common prior about the state of the world, and their conjectures must be common knowledge. Common knowledge of rationality is not one of the necessary conditions for the Nash equilibrium to obtain. In standard game theory models, players' conjectures are not assumed to be mutually (commonly) known. If players' conjectures are private and uncorrelated, they may choose best-response strategies to their private beliefs, and the combination of their best response strategies may not be a Nash equilibrium of the game. For a detailed discussion and proofs, see Aumann and Brandenburger 1995 and Perea 2012.

out the strategy that s/he has to choose in order to make the attainment of that outcome possible. This theory offers an explanation of how people may use a certain concept of mutual advantage to resolve certain non-cooperative games with multiple rationalizable outcomes. However, it achieves that by making a potentially conceptually problematic departure from the orthodox notion of individual rationality – one of the central principles of the orthodox game theory⁵.

The theory of virtual bargaining (Misyak and Chater 2014 and Misyak et al. 2014) is a hypothetical, or fictitious, bargaining model. It is an attempt to incorporate the notion of mutual advantage into game theoretic analysis of non-cooperative games in a way compatible with the basic principles of orthodox game theory. The theory suggests that decision-makers choose their strategies on the basis of what strategy profile(s) they would agree to play if they could openly bargain – engage in real negotiations, in which each player can communicate his or her offers to the other players and receive their counteroffers.

An idea that certain principles of bargaining theory can be used in the game theoretic analysis of other types of non-cooperative games is not entirely new: Conceptual connections between bargaining and the equilibrium (rationalizable strategy) selection problems in non-cooperative games have been discussed, among others, by Raiffa (1953), Luce and Raiffa (1957), Aumann (1959), Schelling (1960), Myerson (1991), Moreno and Wooders (1996), and Ambrus (2006, 2009). The virtual bargaining theory is novel in a sense that the suggested model of hypothetical bargaining is treated not merely as a descriptive tool providing approximately accurate descriptions of people's choices (the so-called *as if* model), but as a tool providing an approximately accurate description of the outcome of an actual process of mental simulation of open bargaining by which people arrive at their strategy choices in non-cooperative games. In other words, the proponents of the theory claim that people actually engage in mental, or 'virtual', simulation of the open bargaining process in order to resolve non-cooperative games (see Misyak and Chater 2014, Misyak et al. 2014).

Hypothetical bargaining is a conceptually appealing idea. In bargaining games where players' agreements are not binding, the set of feasible agreements is the set of correlated equilibria. A bargaining solution is a correlated equilibrium which satisfies a number of intuitively desirable properties. According to Myerson, a bargaining solution can be interpreted as a *reasonably accurate expectation* of the outcome of open bargaining process involving self-oriented individuals of roughly equal bargaining abilities (see Myerson

⁵For a detailed discussion, see chapter 2.

1991). Therefore, it seems reasonable to believe that certain properties of bargaining solutions that decision-makers recognize as being desirable may also play a role in players' identification of mutually beneficial solutions of other types of non-cooperative games.

In addition, bargaining theory is a branch of non-cooperative game theory: It relies on the same basic principles of the orthodox game theory as solution concepts of non-cooperative games (for extensive discussion, see Myerson 1991). Unlike a team-reasoning decision-maker whose aim is to maximally advance the interests of a team, a bargainer is a self-oriented decision-maker – an individual who aims to maximally advance his or her personal interests, and only cares about the interests of other interacting individuals insofar as their actions may promote or hinder the advancement of his or her own personal interests. Like a best-response reasoner, a bargainer is assumed to deviate from the agreement in situations where such deviation is personally beneficial. For this reason, bargaining solutions have some conceptually appealing stability properties.

Finally, like the team reasoning theory, the theory of hypothetical bargaining suggests that players aim to resolve games by identifying and implementing a *mutually advantageous* solution. This hypothesis has some indirect empirical support. Colman and Stirk (1998) conducted an experiment with coordination games, in which participants were asked to report the reasons of their strategy choices. The results suggest that a substantial proportion of people use some notion of mutual advantage when reasoning about non-cooperative games⁶.

However, the theory of hypothetical bargaining is relatively new and therefore has substantial conceptual limitations. A formal model representing the process of mental bargaining is not yet available (for discussion, see chapter 1). One of the fundamental questions pertaining to the theory of hypothetical bargaining which, I believe, does not have a satisfactory answer is what properties a strategy profile must have in order to be identified by hypothetical bargainers as the hypothetical bargaining solution of a game. Misyak and Chater suggest that the 'existing formal accounts of explicit bargaining, such as Nash's theory of bargaining, while incomplete, are nonetheless useful as a starting point for the analysis of virtual bargaining' (Misyak and Chater 2014: 4). In other words, at its current state the the-

⁶Colman and Stirk (1998) report results from multiple common interest and mixed motive games, including the Hi-Lo game, the Chicken game, the Stag Hunt Game, the Battle of the Sexes game, the Prisoner's Dilemma game and the Deadlock game. Their results suggest that a substantial proportion of participants justified their choices in one-shot games by appealing to some notion of mutual benefit ('most points for both', 'mutual benefit', *etc.*)

ory is, essentially, the Nash bargaining solution applied to non-cooperative games.

In this chapter I will argue that the application of the Nash bargaining solution to non-cooperative games is problematic for two reasons. First, the Nash bargaining solution function is not sensitive to the relevant information about the possible *alternative* allocations of players' personal payoff gains, and therefore it does not offer a compelling answer to the question of how hypothetical bargainers identify a hypothetical bargaining solution in games with multiple (weakly) Pareto optimal alternatives, each of which is associated with a different allocation of players' personal payoff gains.

Second, I will argue that the standard Nash bargaining solution can only be applied to cases where players have information about each other's preferences over the lotteries over the set of possible agreements (i.e. cardinal preferences over the possible agreements), and so cannot be meaningfully applied to cases where only ordinal information about preferences is available to the interacting decision-makers.

In this chapter I propose a *benefit-equilibrating* (later abbreviated as BE) hypothetical bargaining solution concept for non-cooperative games, which is broadly in line with the principles underlying Conley and Wilkie's (2012) ordinal egalitarian solution for Pareto-optimal point selection problems with finite choice sets. I will argue that the proposed solution concept can be applied to cases where players only have information about each other's ordinal payoffs, as well as to cases where interpersonal comparisons of decision-makers' cardinal payoffs are assumed not to be meaningful. I offer both the ordinal and the cardinal versions of this solution concept, discuss their formal properties, and illustrate their application in the formal analysis of non-cooperative games with a number of experimentally relevant examples.

The rest of the chapter is structured as follows. In section 2 I discuss the virtual bargaining theory and the reasons of why the use of the Nash bargaining solution for representation of the outcomes of hypothetical bargaining in non-cooperative games may be conceptually problematic. In sections 3 and 4 I propose the ordinal and the cardinal versions of the benefit-equilibrating (BE) bargaining solution for two player games. In section 5 I discuss a version of the BE solution concept for n -player games. With section 6 I conclude and discuss some of the limitations of the proposed model.

3.2 Hypothetical Bargaining

3.2.1 Misyak and Chater’s Virtual Bargaining Model

According to virtual bargaining theory, hypothetical bargainers resolve non-cooperative games by identifying those strategy profiles which, they believe, they would agree to play if they could openly bargain – engage in real negotiations, in which each player communicates his or her offers to the other players and receives their counteroffers. A player who reasons as a hypothetical bargainer interprets all the pure and mixed strategy profiles of a game as *possible agreements*. S/he then identifies a set of *feasible agreements* – a subset of possible agreements, where each element of this subset is a strategy profile, such that no player can exploit the other players by deviating from it. The player then identifies a feasible agreement (or agreements) which, s/he believes, the players would agree to play in open bargaining, and plays his or her part in realizing that agreement, provided that s/he has a reason to believe that the other players are hypothetical bargainers and will carry out their end of that agreement by choosing the appropriate strategies. An agreement (or agreements) identified as the hypothetical bargaining solution is the mutually beneficial and agreeable solution of the game for players who reason as hypothetical bargainers (for details, see Misyak and Chater 2014 and Misyak et al. 2014).

To grasp the intuition behind this model of reasoning, consider the Hi-Lo game depicted in Figure 3.2:

	<i>hi</i>	<i>lo</i>
<i>hi</i>	2, 2	0, 0
<i>lo</i>	0, 0	1, 1

Figure 3.2: Hi-Lo game

This game has two Nash equilibria in pure strategies – (hi, hi) and (lo, lo) . The third rational solution is a mixed Nash equilibrium $(\frac{1}{3}hi, \frac{2}{3}lo; \frac{1}{3}hi, \frac{2}{3}lo)$. The best-response reasoners who cannot communicate with each other face a coordination problem. Yet if the players were to negotiate a joint action plan, they would immediately agree to play (hi, hi) , since this strategy profile is associated with the best possible payoffs for both players. The joint action plan to realize strategy profile (hi, hi) is a self-enforcing, or feasible, agreement: If one player chooses *hi*, the self-oriented co-player can do nothing more personally advantageous than to choose *hi* as well, since this strategy

is a unique best response to opponent's choice of strategy hi .

The fundamental question pertaining to this theory is what formal bargaining solution concept could be used to identify the outcomes that hypothetical bargainers would recognize as hypothetical bargaining solutions of non-cooperative games. Misyak and Chater suggest that the 'goodness of a feasible bargain is, following Nash's theory of bargaining, the product of the utility gains to each player (relative to a no-agreement baseline) of adhering to that agreement' (Misyak and Chater 2014: 4).

Misyak and Chater do not provide a formal model of how the Nash bargaining solution is supposed to be applied to non-cooperative games, yet an attempt to (re)construct it will be made in order to illuminate some of the non-trivial technical assumptions of this theory.

Suppose that two hypothetical bargainers are playing a normal form game Γ in which each player $i \in \{1, 2\}$ has a set of pure strategies S_i . A mixed strategy of player $i \in \{1, 2\}$ is a probability distribution over S_i . Let Σ_i be a set of all such probability distributions and $\sigma_i \in \Sigma_i$ be a mixed strategy of $i \in \{1, 2\}$, where $\sigma_i(s_i)$ is a probability assigned to pure strategy $s_i \in S_i$. A mixed strategy outcome is defined as a mixed strategy profile $\sigma = (\sigma_1, \sigma_2)$.

Let $\Sigma = \Sigma_1 \times \Sigma_2$ be the set of all the mixed strategy profiles of Γ . According to the principles of Nash bargaining theory, each player's preferences over the possible agreements must capture their attitude to risk (for an extensive discussion, see Nash 1950a). Each player's preferences must therefore be defined over the set of lotteries. Each lottery 'prize' is a particular combination of players' mixed or pure strategies, or strategy profile.

Let $\mathcal{L}(\Sigma)$ be a set of lotteries over Σ and $u_i : \mathcal{L}(\Sigma) \rightarrow \mathbb{R}$ be a payoff function of player $i \in \{1, 2\}$. From the set of possible agreements, hypothetical bargainers must establish a disagreement point – a utility pair $d = (u_1^{ref}, u_2^{ref})$ representing each player's expectation of a personal outcome that s/he would get if the players were to fail to reach an agreement. A set of *possible* agreements \mathcal{B} can then be defined as a set of utility pairs where each player gets at least his or her disagreement payoff:

$$\mathcal{B} = \{(u_1(\sigma), u_2(\sigma)) : u_i(\sigma) \geq u_i^{ref} \forall i \in \{1, 2\}\} \quad (3.1)$$

In standard bargaining games the set of self-enforcing agreements in cases where no external enforcement is available is taken to be the set of *correlated equilibria* of a game⁷. If the players were able to communicate, they could,

⁷In cases where players' agreements are enforced by an external party, the feasibility set includes all the *possible* agreements. If the external enforcer is not available, the set of feasible agreements is composed only of correlated equilibria – the self-enforcing agreements of the game. Misyak and Chater seem to suggest that virtual bargaining mimics the procedure of bargaining where external an external enforcer is not available.

in principle, implement any correlated equilibrium. In the context of non-cooperative games, however, the players cannot communicate. They cannot agree to coordinate their actions by observing some correlation device, such as, for example, a toss of a fair coin. Therefore, it seems natural to assume that only the Nash equilibria⁸ are the *feasible* (i.e. both implementable and self-enforcing) agreements in a non-cooperative game. A set of feasible agreements $\mathcal{F} \in \mathcal{P}(\mathcal{B})$ can thus be defined in the following way:

$$\mathcal{F} = \{(u_1(\sigma), u_2(\sigma)) \in \mathcal{B} : \sigma \in \Sigma^{NE}\}. \quad (3.2)$$

The Nash bargaining solution function $\mathcal{N}(\cdot, \cdot)$ satisfies, for every (\mathcal{F}, d) ,

$$\mathcal{N}(\mathcal{F}, d) \in \operatorname{argmax}_{(u_1(\sigma), u_2(\sigma)) \in \mathcal{F}} (u_1(\sigma) - u_1^{ref})(u_2(\sigma) - u_2^{ref}). \quad (3.3)$$

3.2.2 The Limitations of the Model

The standard bargaining solution concepts, including the Nash bargaining solution, have been developed for a specific class of games, known as bargaining problems. In a standard bargaining problem, there exists a unique disagreement point – an outcome that obtains when individuals fail to reach an agreement. Each player can enforce the disagreement outcome: If s/he decides not to accept any offers, the bargainers end up with the disagreement outcome. For example, in the Nash (1950a) bargaining problem a disagreement point is assumed to be an outcome in which both players gain nothing. In axiomatic bargaining theory, players' disagreement points are used to determine each player's personal utility gains from each feasible agreement, and thus it plays a fundamental role in formal characterizations of the standard bargaining solutions⁹. In strategic (alternating offers) bargaining models, disagreement points are interpreted as threat points: At each step of the bargaining process, each player has the ability to reject the opponent's offer and force him or her to consider a counteroffer by threatening to play his or her disagreement strategy, which would harm the opponent by bringing him

⁸Misyak and Chater (2014) suggest that real-world decision-makers may be using a less restrictive feasibility criterion than the one which underlies the Nash equilibrium concept. That criterion will be discussed separately in section 3.5.2.

⁹In axiomatic bargaining theory, the disagreement point is used to identify the feasible agreements which satisfy a certain set of desirable properties, such as Pareto efficiency, symmetry, independence of irrelevant alternatives, proportionality, *etc.* See, for example, Luce and Raiffa 1957, Kalai and Smorodinsky 1975, Kalai 1977, Myerson 1977, Roth 1979.

or her down to his or her personal disagreement payoff as well¹⁰.

The Nash bargaining solution has been developed to resolve a specific type of game, known as the Nash bargaining problem. In the standard formulation of the Nash bargaining problem, two players have to decide on how to split a perfectly divisible good. Each player's utility function represents his or her preferences over lotteries over the set of feasible allocations of the good. The Nash bargaining solution of this problem is a *unique* distribution of the good¹¹.

In other types of non-cooperative games, however, player's utility function is not always defined over allocations of some divisible good. They may represent any motivations which are relevant for player's evaluation of the possible outcomes. A non-cooperative game may have multiple feasible agreements which maximize the Nash product, yet each agreement may be associated with a *different allocation* of personal payoff gains. Since hypothetical bargainers are assumed to be self-oriented decision-makers, it stands to reason to assume that they would not be indifferent between agreements associated with different personal payoff gains, and so the question of how a conflict over allocations of players' personal payoff gains would be resolved becomes a crucial one.

For example, consider the two player four strategy coordination game with four weakly Pareto optimal outcomes¹² depicted in Figure 3.3(a). To simplify the analysis, it will be assumed that players only consider pure strategy outcomes as possible agreements. There are three weakly Pareto optimal pure strategy Nash equilibria in this game: $(s1, t1)$, $(s2, t2)$ and $(s3, t3)$. Notice that each player can guarantee himself or herself a minimum payoff of 1 by playing strategy $s4/t4$, which is player's pure *maximin* strategy. A profile of players' *maximin* strategies $(s4, t4)$ is a Pareto inefficient Nash equilibrium.

Suppose that in case of disagreement each player reverts to playing his or her *maximin* strategy $s4/t4$, and so the disagreement point is the Nash equilibrium $(s4, t4)$. Relative to this disagreement point, the Nash bargaining solution is the Nash equilibrium $(s1, t1)$, since it maximizes the product of

¹⁰In a strategic bargaining model with exogenous risk of breakdown, there is an additional assumption that bargaining will terminate without agreement, with players getting their disagreement payoffs. For an extensive discussion of strategic bargaining models, see Binmore 1980, Rubinstein 1982 and Binmore et al. 1986.

¹¹For an extensive discussion of the Nash bargaining theory, see Nash 1950a, Luce and Raiffa 1957, and Myerson 1991.

¹²An allocation of payoffs associated with an outcome is said to be weakly Pareto optimal if there is no alternative outcome associated with an allocation of payoffs which makes each interacting player strictly better off.

	$t1$	$t2$	$t3$	$t4$
$s1$	10, 9	0, 0	0, 0	0, 1
$s2$	0, 0	9, 9	0, 0	0, 1
$s3$	0, 0	0, 0	9, 9	0, 1
$s4$	1, 0	1, 0	1, 0	1, 1

(a)

	$t1$	$t2$	$t3$	$t4$
$s1$	10, 9	0, 0	0, 0	0, 1
$s2$	0, 0	9, 9	0, 0	0, 1
$s3$	0, 0	0, 0	9, 10	0, 1
$s4$	1, 0	1, 0	1, 0	1, 1

(b)

	$t1$	$t2$	$t3$	$t4$
$s1$	10, 2	0, 0	0, 0	0, 1
$s2$	0, 0	4, 4	0, 0	0, 1
$s3$	0, 0	0, 0	2, 10	0, 1
$s4$	1, 0	1, 0	1, 0	1, 1

(c)

Figure 3.3: 4x4 games with three weakly Pareto optimal Nash equilibria

players' payoff gains. Notice that this outcome is associated with the highest possible payoffs for both players. Given the set of weakly Pareto optimal Nash equilibria available in this game, no player could raise an objection against an offer to play strategy profile $(s1, t1)$.

This, however, would not be the case in games depicted in Figures 3.3(b) and 3.3(c). Relative to disagreement point $(s4, t4)$, the game depicted in Figure 3.3(b) has two Nash bargaining solutions — $(s1, t1)$ and $(s3, t3)$. Notice, however, that players have conflicting preferences over the two solutions: The row player's most preferred outcome is $(s1, t1)$, while the column player's most preferred outcome is $(s3, t3)$. This means that players' personal payoffs could not be simultaneously maximized: *If one of the outcomes were chosen for implementation, one of the players would maximize his or her personal payoff, while the other player would lose an opportunity to maximally advance his or her personal interests and, relative to the maximum payoff attainable in this game, suffer a certain payoff loss.* Therefore, if one of the two Nash bargaining solutions were chosen for implementation, a disadvantaged player could raise a reasonable objection that the offer is 'unreasonable', since it prioritizes the advancement of one player's personal interests over another's. If, on the other hand, the players were to choose outcome $(s2, t2)$, both of them would suffer a certain loss of maximum attain-

able personal payoff, but neither of them could raise an objection that the advancement of the opponent's personal interests were given priority over the advancement of his or her own personal interests.

A similar conflict of interests would arise in game 3.3(c). In this game, all three weakly Pareto optimal Nash equilibria maximize the product of players' payoffs, relative to disagreement point (s_4, t_4) . Notice that each of the three agreements is associated with a different allocation of players' personal payoff gains. If either outcome (s_1, t_1) or (s_3, t_3) were chosen for implementation, one of the players would maximize his or her personal payoff, while the other player would get a payoff which is only slightly higher than his or her *maximin* payoff. The disadvantaged player could raise a reasonable objection that, *given the set of feasible agreements available in this game*, the offer is 'unreasonable'. If the players were to choose outcome (s_2, t_2) , both of them would suffer a substantial loss of maximum attainable personal payoff. However, neither of them could raise an objection that the advancement of the opponent's personal interests were given priority over the advancement of his or her own personal interests. The Nash bargaining solution, however, does not single out any of the three Nash equilibria as the most reasonable allocation of personal gains. This and the previous example suggest that an arbitration scheme based on the Nash bargaining solution concept may not capture all the considerations which would be relevant for players when resolving games with multiple weakly Pareto optimal feasible agreements¹³.

Each player could threaten the opponent to end the negotiations if s/he were to deem the opponent's offer unreasonable. The disadvantaged player could simply revert to playing his or her *maximin* strategy (s_4/t_4) , thus leaving the opponent no better option than playing his or her *maximin* strategy as well. Alternatively, the disadvantaged player could simply leave the negotiations without giving any indication as to what s/he intends to do (for an extensive discussion of this threat strategy, see Luce and Raiffa (1957)). In such a situation of strategic uncertainty, the opponent could either play his or her *maximin* strategy, or attempt to guess the disadvantaged player's

¹³Misyak and Chater (2014) discuss a version of the Battle of the Sexes game, in which one Nash equilibrium is associated with what they refer to as 'asymmetric payoffs' (1, 11), while the other Nash equilibrium is associated with what they deem to be 'mutually good' payoffs (10, 9). They suggest that the Nash equilibrium with 'mutually good' payoffs is a more likely bargaining outcome than the Nash equilibrium with 'asymmetric payoffs', since the disadvantaged player will likely reject the offer with 'asymmetric payoffs'. However, they do not consider a situation where multiple outcomes which maximize the Nash product (in their example, the mutually good outcome is the unique Nash bargaining solution of the game) may be associated with different allocations of payoff gains, nor do they offer a theoretical explanation of how a choice between multiple Nash bargaining solutions could be made without interpersonal comparisons of payoffs.

strategy. In either case, the opponent would risk getting a strictly lower payoff than the one that s/he would attain if the players were to agree on playing any of the three weakly Pareto optimal Nash equilibria. Notice that disadvantaged player's threat to end the negotiations in response to what s/he perceives as an unreasonable offer *satisfies Raiffa's credible threat condition*, since, in case of failed negotiations, the disadvantaged player would, *relative to payoffs associated with an unreasonable offer*, face a risk of suffering a smaller payoff loss than the player making an unreasonable offer (for extensive discussion, see Raiffa 1953 and Luce and Raiffa 1957).

Because of the risks associated with a failure to reach an agreement, a rational player should be motivated to reach an agreement rather than to face the consequences of failed negotiations. Therefore, each player should be motivated to focus on the feasible solutions which would minimize the risk of failed negotiations — the set of feasible agreements which the opponent would deem reasonable.

There are several game theoretic models which purport to explain players' choices in experimental games by incorporating fairness considerations and other types of pro-social preferences into players' payoff functions. One of the well-known models is the inequity aversion theory suggested by Fehr and Schmidt (1999). These theories, although useful in explaining players' choices in games with material payoffs, cannot be applied to games where players' payoffs are their von Neumann and Morgenstern utilities, which are supposed to represent all the motivations relevant for players' evaluations of outcomes, including, among other things, players' pro-social preferences, such as inequity aversion, altruism, sensitivity to social norms, and so on.

Another limitation of the standard Nash bargaining solution is that it can only be applied to cases where players' cardinal preferences over lotteries are common knowledge among the interacting players. This is a strong epistemic requirement: In many real-world interdependent decision problems, people, at best, know each other's ordinal preferential rankings of feasible alternatives. The standard Nash bargaining solution cannot be applied to such cases. However, it is obvious that people engage in negotiations even in situations where they only have ordinal information about each other's preferences. If hypothetical bargaining is supposed to represent the actual process of reasoning by which real-world decision-makers resolve games, then solution concept's applicability to cases where decision-makers have rudimentary information about each other's preferences is a desirable feature.

In the following sections, I will suggest an alternative bargaining-based explanation of how players resolve the payoff allocation problems. I will suggest two versions of the formal benefit-equilibrating (BE) solution concept

– one for ordinal and one for cardinal cases. I will argue that a certain type of comparison of foregone opportunities plays an important role in hypothetical bargaining, and that the benefit-equilibrating solution offers a plausible explanation of how such comparisons of foregone opportunities may determine players' choices in non-cooperative games.

3.3 The Ordinal Benefit-Equilibrating Solution

3.3.1 The Intuition Behind the Ordinal BE Solution

In every negotiations, each self-oriented individual wants to maximally advance his or her personal interests. S/he is therefore interested in pushing the other bargaining party or parties to accept as many of his or her initial demands as possible. In cases where individuals' personal interests coincide perfectly, they may reach an agreement without giving up any of their initial demands. However, in cases where individuals have conflicting interests (e.g. in cases where their demands cannot be met simultaneously), an agreement can only be reached by at least one of the bargaining parties making a *concession* – giving up some of the initial demands. In such situations, a self-oriented negotiator will be interested in reaching an agreement which minimizes the number of his or her foregone initial demands, since such an agreement would maximize the advancement of negotiator's personal interests. A negotiator can therefore evaluate the 'goodness' of each feasible agreement on the basis of the number of initial demands that s/he would have to forego in order for that agreement to be reached: An agreement which could be reached with a smaller number of foregone initial demands should always be deemed better than the one which would require a larger sacrifice of initial demands (Zhang and Zhang 2008).

In addition, the bargainers can use another criterion for evaluating the feasible bargaining agreements. Assuming that each bargainer knows the set of each opponent's initial demands, s/he can also determine the number of initial demands that each of the bargainers would have to forego in order to reach a particular agreement. Each bargainer can therefore compare the number of initial demands that s/he would have to give up in order to reach a particular agreement with the number of initial demands that would have to be sacrificed by every other bargainer. The 'reasonableness' of each feasible agreement can then be evaluated on the basis of the distributions of foregone initial demands.

The ordinal benefit-equilibrating (BE) solution is based on the principle that hypothetical bargainers compare the acceptability of each feasible

agreement by comparing the distributions of foregone initial demands associated with each feasible agreement: An agreement with a more equitable distribution of foregone initial demands is deemed more acceptable than the one with a less equitable distribution of foregone initial demands. In other words, it is assumed that hypothetical bargainers not only care about the properties of the agreement itself, but also about how that agreement is reached. Each feasible agreement can be reached by each bargainer giving up a certain number of personal initial demands. An agreement associated with a more equitable distribution of foregone initial demands among the interacting individuals gets picked over an agreement with a less equitable distribution of foregone initial demands.

These principles can be applied to the analysis of non-cooperative games where players only have ordinal information about each other's preferences over possible outcomes. For example, consider a simple two player three strategy non-cooperative ordinal coordination game depicted in Figure 3.4. The left and the right number in each cell represent the row and the column player's *ordinal preferences over outcomes* respectively.

	$t1$	$t2$	$t3$	$t4$
$s1$	100, 3	0, 0	0, 0	0, 0
$s2$	0, 0	60, 5	0, 0	0, 0
$s3$	0, 0	0, 0	40, 9	0, 0
$s4$	0, 0	0, 0	0, 0	20, 1

Figure 3.4: Ordinal coordination game

Suppose that players do not know each other's preferences over lotteries over the pure strategy profiles of the game, and so only consider pure strategy profiles as possible agreements. It will be assumed that the set of feasible (i.e. self-enforcing) agreements for such players is the set of pure strategy Nash equilibria. The game depicted in Figure 3.4 has four Nash equilibria in pure strategies – $(s1, t1)$, $(s2, t2)$, $(s3, t3)$ and $(s4, t4)$. Each feasible agreement can be interpreted as a possible state of the world that players could bring about if they were to agree on playing a particular combination of strategies. Assuming that players' ordinal preferences over outcomes of this game are common knowledge, each player could determine the number of *preferred alternative agreements* that each player would forego if each of the feasible agreements were chosen to be implemented. If the outcome $(s1, t1)$ were chosen, the row player's personal interests would be maximally

advanced, since s/he prefers this agreement over all the other feasible alternative agreements (in this case over all the possible outcomes) of this game. The row player would therefore *lose no opportunities to advance his or her personal interests*. The column player, on the other hand, prefers the outcome (s_2, t_2) and the outcome (s_3, t_3) over the outcome (s_1, t_1) . If the outcome (s_1, t_1) were chosen, the column player would forego two preferred alternative agreements.

If the outcome (s_3, t_3) were chosen, the column player would forego no opportunities to advance his or her personal interests, since s/he prefers this agreement over all the other feasible alternative agreements of this game. The row player, however, prefers the outcome (s_1, t_1) and the outcome (s_2, t_2) over the outcome (s_3, t_3) . If the outcome (s_3, t_3) were chosen, s/he would forego two preferred alternative agreements.

By choosing outcome (s_2, t_2) , each player would forego one preferred alternative agreement: The row player prefers the outcome (s_1, t_1) over the outcome (s_2, t_2) , while the column player prefers the outcome (s_3, t_3) over the outcome (s_2, t_2) . If the outcome (s_4, t_4) were chosen, each player would lose three preferred alternative agreements: Each player prefers outcomes (s_1, t_1) , (s_2, t_2) and (s_3, t_3) over the outcome (s_4, t_4) .

Assuming that each player is self-oriented, s/he will always prefer a feasible agreement associated with a smaller number of foregone preferred alternatives over a feasible agreement associated with a larger number of foregone preferred alternatives. If this preference relation is common knowledge, each player must be able to construct, for every player of the game, a preferential ranking of feasible agreements based on the numbers of foregone preferred alternatives. These rankings are shown in Table 3.1:

	Agreement	Foregone alt.		Agreement	Foregone alt.
Row:	(s_1, t_1)	0	Column:	(s_3, t_3)	0
	(s_2, t_2)	1		(s_2, t_2)	1
	(s_3, t_3)	2		(s_1, t_1)	2
	(s_4, t_4)	3		(s_4, t_4)	3

Table 3.1: Players' foregone preferred alternatives

Both bargainers would easily agree to *restrict* their negotiations to a subset of feasible agreements including outcomes (s_1, t_1) , (s_2, t_2) and (s_3, t_3) . Such a restriction of the bargaining set is clearly mutually beneficial: For each bargainer, any agreement in the aforementioned subset guarantees a strictly lower number of foregone preferred alternatives than the agreement (s_4, t_4) . Therefore, each player should strictly prefer any agreement within the aforementioned subset over the agreement (s_4, t_4) . In other words, self-

oriented players would easily agree not to consider weakly Pareto dominated feasible agreements. Among the feasible agreements (s_1, t_1) , (s_2, t_2) and (s_3, t_3) , however, there is no mutually beneficial agreement: Each of the agreements in the subset is associated with a particular distribution of foregone preferred alternatives, yet no agreement in this subset is, relative to any other agreement in that subset, associated with strictly lower numbers of foregone preferred alternatives for both players. It means that there is no agreement in the subset which, relative to any other agreement in that subset, would make both players strictly better off. In this case, hypothetical bargainers could evaluate the feasible agreements by comparing how the foregone preferred alternatives would be distributed among the interacting individuals if each of the agreements were chosen.

Since players' personal interests cannot be maximally advanced simultaneously, they would have to agree on how to distribute the losses of the preferred alternative agreements: If either the agreement (s_1, t_1) or the agreement (s_3, t_3) were chosen, one of the players would forego none of the preferred alternative agreements, while the other one would forego two. If the outcome (s_2, t_2) were chosen, each player would forego one preferred alternative agreement.

Notice that each player could threaten the opponent to end the negotiations if s/he were to deem the opponent's offer unreasonable. The disadvantaged player's threat to end the negotiations in response to an unreasonable offer satisfies Raiffa's credible threat condition, since, in case of failed negotiations, the disadvantaged player would face a risk of losing a strictly lower number of foregone preferred alternatives than the player making an unreasonable offer. If the players were to fail to reach an agreement, they would end up with no joint plan on how to resolve the game. In such a situation, each player could revert to playing his or her *maximin* strategy in order to secure the best possible outcome that s/he can guarantee to himself or herself, irrespective of what strategy the other player chooses. However, the *maximin* outcome for each player is worse than agreements (s_1, t_1) , (s_2, t_2) and (s_3, t_3) . Alternatively, each player could attempt to guess the opponent's strategy choice. However, by doing this the player could end up with the worst possible outcome. Because of the risks associated with a failure to reach an agreement, a rational player should be motivated to reach an agreement rather than to face the consequences of failed negotiations.

Notice that outcome (s_2, t_2) is a weakly Pareto optimal agreement which minimizes the difference between the numbers of players' foregone preferred alternatives. In other words, among the three weakly Pareto optimal agreements (s_1, t_1) , (s_2, t_2) and (s_3, t_3) , agreement (s_2, t_2) ensures a maximally

equitable distribution of foregone preferred alternatives. If the agreement (s_2, t_2) were chosen, no player could raise an objection that, given the possible ways in which each player's personal interests could be advanced in this game, the advancement of the other player's personal interests – the minimization of the number of other player's foregone alternatives – has been given priority over the advancement of his or her own personal interests. The pure strategy Nash equilibrium (s_2, t_2) is the *benefit-equilibrating* (BE) *solution* of this game.

The ordinal BE solution concept is based on the principles which are quite similar to the ones underlying Conley and Wilkie's (2012) *ordinal egalitarian bargaining solution* for finite sets of Pareto optimal points. In Conley and Wilkie's (2012) model, each bargainer has an ordinal ranking of Pareto optimal points, based on calculations of player's *cardinalities of the preferred sets of alternatives* of each Pareto optimal point. A cardinality of the preferred set of any Pareto optimal point x is simply the number of Pareto optimal points that a particular bargainer prefers over the point x . An ordinal egalitarian bargaining solution is a Pareto optimal point which, for both players, is associated with equal numbers of foregone preferred alternatives (for an in-depth discussion, axiomatic characterization and proofs, see Conley and Wilkie 2012).

The ordinal BE solution, however, is based on a weaker equity requirement: It is any weakly Pareto optimal outcome which, given a particular set of weakly Pareto optimal outcomes, *minimizes the difference* between the numbers of players' foregone preferred alternatives. This means that, in some games, a BE solution may not be ordinally egalitarian in the sense suggested by Conley and Wilkie. It is, however, a maximally ordinally equitable outcome available in a particular set of feasible weakly Pareto optimal agreements. This weaker equity requirement is based on assumption that hypothetical bargainers would not revert to playing a Pareto suboptimal agreement in games where a strictly ordinally egalitarian and weakly Pareto optimal agreements were not available, but would rather agree on playing a weakly Pareto optimal agreement associated with a maximally equitable distribution of foregone preferred alternatives. Despite this difference, a BE solution, like the ordinally egalitarian solution, singles out solutions which are compatible with an intuitively compelling principle that, in every bargaining problem where bargainers have conflicting preferences over a set of weakly Pareto optimal outcomes, an agreement which, compared to other feasible weakly Pareto optimal agreements, *decreases* the number of foregone preferred alternatives of one player at the expense of *increasing* the number of other player's foregone preferred alternatives will not be agreed upon by

self-oriented bargainers.

3.3.2 Formalization

Let $\Gamma^o = (\{1, 2\}, \{S_i, \succeq_i\}_{i \in \{1, 2\}})$ be any two player ordinal game, in which each player $i \in \{1, 2\}$ has a finite set of strategies S_i . Let $\mathbf{S} = S_1 \times S_2$ denote the set of all pure strategy profiles, or outcomes, of Γ^o . Each player $i \in \{1, 2\}$ has a complete and transitive preference ranking \succeq_i over the set \mathbf{S} . Every pure strategy outcome is a pure strategy profile $\mathbf{s} = (s_1, s_2)$. It will be assumed that players do not know each other's preferences over lotteries over pure strategy outcomes, and only consider pure strategy outcomes as possible agreements.

Let $S_i^{br} \subseteq S_i$ denote the set of rationalizable strategies of $i \in \{1, 2\}$. A strategy $s_i \in S_i$ of $i \in \{1, 2\}$ is rationalizable if, for some opponent j 's strategy $s_{j \neq i} \in S_j$,

$$(s_i, s_j) \succeq_i (\tilde{s}_i, s_j) \quad \forall \tilde{s}_i \in S_i. \quad (3.4)$$

Let $\mathbf{S}^{br} = S_1^{br} \times S_2^{br}$ denote the set of rationalizable strategy profiles of Γ^o . For each rationalizable outcome $\mathbf{s} \in \mathbf{S}^{br}$, we can define the cardinality of the preferred set of alternatives for every player $i \in \{1, 2\}$:

$$C_i(\mathbf{s}, \mathbf{S}^{br}) \equiv \{ |T| \text{ where } \mathbf{s}' \in T \text{ if and only if } \mathbf{s}' \in \mathbf{S}^{br} \text{ and } \mathbf{s}' \succ_i \mathbf{s} \}. \quad (3.5)$$

Let $C_i^{ref}(\mathbf{S}^{br})$ denote the reference point of player $i \in \{1, 2\}$. With respect to the reference point, two definitions seem reasonable. One possibility is to define the reference point as the worst possible outcome in rationalizable strategies. In terms of cardinalities of the preferred sets of alternatives, this reference point can be defined as follows:

$$C_i^{ref}(\mathbf{S}^{br}) = \max_{\mathbf{s} \in \mathbf{S}^{br}} \{ C_i(\mathbf{s}, \mathbf{S}^{br}) \}. \quad (3.6)$$

The intuition behind this definition is as follows: Hypothetical bargainers who fail to reach an agreement in open bargaining have no joint plan on how to play the game. In such a situation of strategic uncertainty, the players may attempt to coordinate their actions by guessing each other's strategy choice (this definition relies on assumption that players cannot choose not to play the game in case of a failure to reach an agreement). If rationality is common knowledge, the players should only consider rationalizable strategies. This means that players should expect any outcome of their attempt to coordinate actions to be a profile of rationalizable strategies. Each player's reference point is the worst possible personal outcome that such an attempt to coordinate actions may yield.

Another possibility is to define the reference point as the *maximin* outcome in rationalizable strategies. In terms of cardinalities of preferred sets of alternatives, the *maximin* outcome in rationalizable strategies can be defined as follows:

$$C_i^{ref}(\mathbf{S}^{br}) = \min_{s_i \in S_i^{br}} \left\{ \max_{s_{j \neq i} \in S_j^{br}} C_i(\mathbf{s}, \mathbf{S}^{br}) \right\}. \quad (3.7)$$

The intuition behind this definition is as follows: If hypothetical bargainers fail to reach an agreement, each decision-maker expects the opponent to choose one of his or her rationalizable strategies, yet s/he is uncertain as to which one of the rationalizable strategies the opponent is going to choose. The player will respond to strategic uncertainty by choosing a strategy which guarantees the best possible outcome (i.e. the outcome with the smallest cardinality of the preferred set of alternatives), irrespective of which one of the rationalizable strategies the opponent is going to choose.

The question of which one of the suggested reference points is the best approximation to how real-world hypothetical bargainers reason about their options in games cannot be answered on the basis of formal theoretical analysis alone. Further empirical research is required to answer this question. It seems reasonable to expect the decision-maker's choice of a reference point to depend on how high his or her personal stakes are in a particular game: It is possible that a decision-maker would adopt a more cautious approach in a game where the personal stakes were high, while would be more willing to risk in a game where the personal stakes were low. *For the purposes of the following theoretical discussion, each hypothetical bargainer's reference point will be assumed to be the worst outcome in rationalizable strategies.* This assumption concerning player's reference point seems reasonable for a model describing hypothetical bargainers' behaviour in experimental games, in which participants' personal stakes are usually relatively low.

Let $\mathcal{A} = \mathcal{P}(\mathbf{S})$ denote the set of *feasible agreements* of Γ^o , where $\mathcal{A} \neq \emptyset$. A pure strategy Nash equilibrium can be defined on the basis of ordinal information about players' preferences. It will be assumed that the set of feasible agreements is the set of pure strategy Nash equilibria of Γ^o :

$$\mathcal{A} = \left\{ \mathbf{s} \in \mathbf{S}^{br} : \mathbf{s} \in \mathbf{S}^{NE} \right\}, \text{ where } \mathbf{S}^{NE} \in \mathcal{P}(\mathbf{S}). \quad (3.8)$$

Since each bargainer is a self-oriented decision-maker, s/he always prefers a feasible agreement associated with a smaller cardinality of the preferred set of alternatives. That is, for any two agreements $\mathbf{s} \in \mathcal{A}$ and $\mathbf{s}' \in \mathcal{A}$ such that $\mathbf{s}' \succ_i \mathbf{s}$, it must be the case that $C_i(\mathbf{s}', \mathcal{A}) < C_i(\mathbf{s}, \mathcal{A})$. When negotiating an implementation of any of the two feasible agreements $\mathbf{s} \in \mathcal{A}$ and $\mathbf{s}' \in \mathcal{A}$, hypothetical bargainers agree to realize agreement $\mathbf{s}' \in \mathcal{A}$ rather than agreement $\mathbf{s} \in \mathcal{A}$ if, for every bargainer $i \in \{1, 2\}$, it is the case that

$C_i(s', \mathcal{A}) < C_i(s, \mathcal{A})$. From this we can define the set of *maximally mutually advantageous* feasible agreements $\mathcal{A}^m \subseteq \mathcal{A}$. A feasible agreement $s \in \mathcal{A}$ is in the set of maximally mutually advantageous feasible agreements $\mathcal{A}^m \subseteq \mathcal{A}$ *only if* there is no alternative feasible agreement $s' \in \mathcal{A}$ which, for both bargainers, is associated with a strictly lower cardinality of the preferred set of alternatives:

$$s \in \mathcal{A}^m \Rightarrow s' \notin \mathcal{A} : C_i(s', \mathbf{S}^{br}) < C_i(s, \mathbf{S}^{br}) \forall i \in \{1, 2\}. \quad (3.9)$$

Let $C_i^{max}(\mathbf{S}^{br})$ denote the cardinality of the preferred set of alternatives associated with the best rationalizable outcome for player $i \in \{1, 2\}$. From the definition of cardinality of the preferred set of alternatives, it follows that the best rationalizable outcome is the one which minimizes i 's cardinality of the preferred set of alternatives:

$$C_i^{max}(\mathbf{S}^{br}) = \min_{s \in \mathbf{S}^{br}} \{C_i(s, \mathbf{S}^{br})\}. \quad (3.10)$$

Let $\varphi_i^o(\cdot, \cdot)$ be an ordinal measure of loss of attainable individual advantage of player $i \in \{1, 2\}$. For an ordinal game, such that, for every $i \in \{1, 2\}$, it is the case that $C_i^{max}(\mathbf{S}^{br}) \leq C_i(s, \mathbf{S}^{br})$ for all $s \in \mathbf{S}^{br}$, a loss of attainable individual advantage of player $i \in \{1, 2\}$ associated with a feasible agreement $s \in \mathcal{A}$ can be defined in the following way:

$$\varphi_i^o(s, \mathbf{S}^{br}) = |(C_i^{max}(\mathbf{S}^{br}) - C_i^{ref}(\mathbf{S}^{br})) - (C_i(s, \mathbf{S}^{br}) - C_i^{ref}(\mathbf{S}^{br}))|, \quad (3.11)$$

which can be simplified to

$$\varphi_i^o(s, \mathbf{S}^{br}) = |C_i^{max}(\mathbf{S}^{br}) - C_i(s, \mathbf{S}^{br})|. \quad (3.12)$$

From definitions (3.5) and (3.10), it follows that $C_i^{max}(\mathbf{S}^{br}) = 0$ for every $i \in \{1, 2\}$. Therefore, for an ordinal game, such that, for every $i \in \{1, 2\}$, it is the case that $C_i^{max}(\mathbf{S}^{br}) \leq C_i(s, \mathbf{S}^{br})$ for all $s \in \mathbf{S}^{br}$, an ordinal measure of loss of attainable individual advantage of $i \in \{1, 2\}$ can be defined in the following way:

$$\varphi_i^o(s, \mathbf{S}^{br}) = |-C_i(s, \mathbf{S}^{br})| = C_i(s, \mathbf{S}^{br}). \quad (3.13)$$

The ordinal BE solution function $\phi^o(\cdot, \cdot)$ satisfies, for every \mathcal{A}^m ,

$$\phi^o(\mathcal{A}, \mathbf{S}^{br}) \in \arg \min_{s \in \mathcal{A}^m} \{|\varphi_i^o(s, \mathbf{S}^{br}) - \varphi_{j \neq i}^o(s, \mathbf{S}^{br})|\}. \quad (3.14)$$

By inserting (3.13) into (3.14), we get the following definition of the ordinal BE solution function:

$$\phi^o(\mathcal{A}, \mathbf{S}^{br}) \in \arg \min_{s \in \mathcal{A}^m} \{|C_i(s, \mathbf{S}^{br}) - C_{j \neq i}(s, \mathbf{S}^{br})|\}. \quad (3.15)$$

3.3.3 Ordinal BE Solution Properties

An ordinal BE solution (not necessarily unique) exists in every finite two player ordinal game with at least one Nash equilibrium in pure strategies, in which each player has a finite set of pure strategies. For any such game, the set $\mathbf{S}^{br} \subseteq \mathbf{S}$ is finite. It follows that $\mathbf{S}^{NE} \in \mathcal{P}(\mathbf{S})$ is also finite. Since $\mathcal{A} = \mathbf{S}^{NE}$, the set \mathcal{A} is finite as well. It is therefore possible to define, for every feasible agreement $\mathbf{s} \in \mathcal{A}$, the cardinality of the preferred set of alternatives for every $i \in \{1, 2\}$. In every finite set $\mathcal{A} = \mathbf{S}^{NE}$, there must exist a feasible agreement $\mathbf{s} \in \mathcal{A}$, such that

$$\mathbf{s}' \notin \mathcal{A} : C_i(\mathbf{s}', \mathbf{S}^{br}) < C_i(\mathbf{s}, \mathbf{S}^{br}) \forall i \in \{1, 2\}. \quad (3.16)$$

It follows that $\mathbf{s} \in \mathcal{A}^m$, which means that $\mathcal{A}^m \neq \emptyset$, and so a BE solution exists.

Every ordinal BE solution is a *weakly Pareto optimal feasible agreement*. Let $\mathcal{A}^{wpo} \subseteq \mathcal{A}$ denote the set of weakly Pareto feasible agreements of Γ^o . A feasible agreement $\mathbf{s} \in \mathcal{A}$ belongs to a set of weakly Pareto optimal feasible agreements *only if* there is no alternative feasible agreement $\mathbf{s}' \in \mathcal{A}$ such that $\mathbf{s}' \succ_i \mathbf{s} \forall i \in \{1, 2\}$. In terms of cardinalities of preferred sets, this condition can be defined as follows:

$$\mathbf{s} \in \mathcal{A}^{wpo} \Rightarrow \mathbf{s}' \notin \mathcal{A} : C_i(\mathbf{s}', \mathbf{S}^{br}) < C_i(\mathbf{s}, \mathbf{S}^{br}) \forall i \in \{1, 2\}. \quad (3.17)$$

Condition (3.17) is equivalent to condition (3.9), which implies that $\mathcal{A}^{wpo} = \mathcal{A}^m$. Since the ordinal BE solution function always picks a subset of $\mathcal{A}^m \subseteq \mathcal{A}$, it follows that BE solution is always a weakly Pareto optimal feasible agreement.

Weak Pareto optimality property is intuitively compelling: It seems unreasonable to expect the intelligent players who understand both the structure of the game and each other's motivations to implement an agreement if there is another feasible agreement which is deemed strictly better by every player. Unlike a strict Pareto improvement, which may not make every interacting individual better-off, a weak Pareto improvement is always beneficial to every interacting individual, and so it seems reasonable to expect it to be chosen by self-oriented decision-makers – individuals who aim to advance their personal interests as much as possible¹⁴.

¹⁴An outcome x is said to be a strict Pareto improvement over the outcome y if it makes at least one individual better-off without making anyone else worse off. A self-oriented individual who does not gain anything from implementing a strict Pareto improvement x has no personal incentive to implement it. An outcome x is said to be a weak Pareto improvement over the outcome y if it makes everyone strictly better off. In case of a weak Pareto improvement, each interacting individual has a personal incentive to implement it.

Every ordinal BE solution is *invariant under additions of weakly Pareto irrelevant alternatives*. In other words, a BE solution will remain invariant under additions of rationalizable strategy profiles which are weakly Pareto dominated by every feasible agreement in the set of maximally mutually advantageous feasible agreements. Suppose that an ordinal game $\Gamma^{o'}$ is derived from an ordinal game Γ^o by extending it with a strategy profile $\mathbf{s}^* \in \mathbf{S}^{br}$. Recall that $\mathcal{A}^m = \mathcal{A}^{wpo}$ for every Γ^o . A strategy profile $\mathbf{s}^* \in \mathbf{S}^{br}$ is a weakly Pareto irrelevant alternative if and only if, for every $i \in \{1, 2\}$, it is the case that $\mathbf{s} \succ_i \mathbf{s}^* \forall \mathbf{s} \in \mathcal{A}^m$. In terms of cardinalities of the preferred sets of alternatives, strategy profile $\mathbf{s}^* \in \mathbf{S}^{br}$ is a weakly Pareto irrelevant alternative if and only if, for every $i \in \{1, 2\}$,

$$C_i(\mathbf{s}^*, \mathbf{S}^{br}) > C_i(\mathbf{s}, \mathbf{S}^{br}) \quad \forall \mathbf{s} \in \mathcal{A}^m. \quad (3.18)$$

It follows that $\mathbf{s}^* \notin \mathcal{A}^m$, which means that ordinal games Γ^o and $\Gamma^{o'}$ are such that $\mathcal{A}^m = \mathcal{A}^{m'}$. From definitions (3.5) and (3.18), it follows that, for every $i \in \{1, 2\}$, it is the case that $C_i(\mathbf{s}, \mathbf{S}^{br'}) = C_i(\mathbf{s}, \mathbf{S}^{br})$ for every $\mathbf{s} \in \mathcal{A}^{m'}$. Since $\mathcal{A}^{m'} = \mathcal{A}^m$, it follows that $\phi^o(\mathcal{A}', \mathbf{S}^{br'}) = \phi^o(\mathcal{A}, \mathbf{S}^{br})$.

Every ordinal BE solution also satisfies a version of the *individual rationality axiom*. In standard *cardinal bargaining models*, a cardinal bargaining solution function satisfies the axiom of individual rationality if and only if, for every bargaining problem, it selects an agreement which yields each player at least his or her disagreement payoff. The intuition behind this requirement is that a rational bargainer should not accept any agreement which provides him or her a payoff lower than a certain threshold. There are multiple suggestions of how this threshold could be defined. One of the suggestions is that a rational bargainer should not accept an agreement which yields him or her a utility lower than his or her *maximin* payoff (for extensive discussion, see Luce and Raiffa 1957, Roth 1977, and Myerson 1991). In game theoretic context, a *maximin* payoff can be defined as the maximum personal payoff that a player can guarantee to himself or herself, irrespective of what the other players choose to do. In ordinal terms, a *maximin* threshold can be defined as a pure strategy outcome that a player prefers over all the other pure strategy outcomes which s/he can guarantee to himself or herself, irrespective of the opponents' strategy choices. For example, in the ordinal coordination game depicted in Figure 3.4, each player's *maximin* outcome is every pure strategy profile associated with the ordinal payoff of 0.

For any strategy profile $\mathbf{s} \in \mathbf{S}$ (including any non-rationalizable strategy profile), the cardinality of the preferred set of alternatives can be defined as follows:

$$C_i(\mathbf{s}, \mathbf{S}) \equiv \{|T'| \text{ where } \mathbf{s}' \in T' \text{ if and only if } \mathbf{s}' \in \mathbf{S} \text{ and } \mathbf{s}' \succ_i \mathbf{s}\}. \quad (3.19)$$

In terms of cardinalities of preferred sets of alternatives, the ordinal *maximin* threshold of $i \in \{1, 2\}$ can be defined as follows:

$$C_i^{mxm}(\mathbf{S}) = \min_{s_i \in S_i} \left\{ \max_{s_{j \neq i} \in S_j} C_i(\mathbf{s}, \mathbf{S}) \right\}. \quad (3.20)$$

The ordinal BE solution function satisfies the individual rationality requirement if and only if, for every $i \in \{1, 2\}$, every BE solution is a strategy profile $\mathbf{s} \in \mathbf{S}$, such that

$$C_i^{mxm}(\mathbf{S}) \geq C_i(\mathbf{s}, \mathcal{A}) \forall \mathbf{s} \in \arg \min_{\mathbf{s} \in \mathcal{A}^m} \left\{ \left| \varphi_i^o(\mathbf{s}, \mathbf{S}^{br}) - \varphi_{j \neq i}^o(\mathbf{s}, \mathbf{S}^{br}) \right| \right\}. \quad (3.21)$$

Since $\mathcal{A} = \mathbf{S}^{NE}$, it follows that every ordinal BE solution is a Nash equilibrium. If a strategy profile $\mathbf{s} \in \mathbf{S}$ is a Nash equilibrium, the preferences of player $i \in \{1, 2\}$ are as follows:

$$(s_i, s_{j \neq i}) \in \mathbf{S}^{NE} \Rightarrow (s_i, s_j) \geq_i (\tilde{s}_i, s_j) \forall \tilde{s}_i \in S_i. \quad (3.22)$$

In terms of cardinalities of preferred sets, property (3.22) can be defined as follows:

$$(s_i, s_{j \neq i}) \in \mathbf{S}^{NE} \Rightarrow C_i((s_i, s_j), \mathbf{S}) \leq C_i((\tilde{s}_i, s_j), \mathbf{S}) \forall \tilde{s}_i \in S_i. \quad (3.23)$$

Notice that the *maximin* strategy $s_i^{mxm} \in S_i$ of each player $i \in \{1, 2\}$ is such that

$$C_i((s_i^{mxm}, s_{j \neq i}), \mathbf{S}) \leq C_i^{mxm}(\mathbf{S}) \forall s_j \in S_j. \quad (3.24)$$

Every Nash equilibrium of Γ^o must have the following property:

$$(s_i, s_{j \neq i}) \in \mathbf{S}^{NE} \Rightarrow (s_i, s_j) \geq_i (s_i^{mxm}, s_j) \forall i \in \{1, 2\}. \quad (3.25)$$

In terms of cardinalities of preferred sets, property (3.25) can be characterized as follows:

$$(s_i, s_{j \neq i}) \in \mathbf{S}^{NE} \Rightarrow C_i(s_i, s_j) \leq C_i(s_i^{mxm}, s_j) \forall i \in \{1, 2\}. \quad (3.26)$$

Since every ordinal BE solution is a Nash equilibrium, the individual rationality requirement is always satisfied.

Every ordinal BE solution is *invariant under additions of non-rationalizable outcomes*. Suppose that $f^o(\cdot)$ is some ordinal bargaining solution function. The axiom requires that, for any two games Γ^o and $\Gamma^{o'}$, such that $\mathbf{S}^{br} = \mathbf{S}^{br'}$, it must be the case that $f^o(\mathbf{S}^{br}) = f^o(\mathbf{S}^{br'})$.

Since C_i^{max} and C_i^{ref} are, for every $i \in \{1, 2\}$, associated with outcomes in $\mathbf{S}^{br} \subseteq \mathbf{S}$, the cardinality of the preferred set of alternatives associated with

every $\mathbf{s} \in \mathbf{S}^{br}$ remains invariant under additions of outcomes which are such that $\mathbf{s} \notin \mathbf{S}^{br}$.

For any two games Γ^o and $\Gamma^{o'}$, such that $\mathbf{S}^{br} = \mathbf{S}^{br'}$, it must be the case that $\mathbf{S}^{NE} = \mathbf{S}^{NE'}$. Since $\mathcal{A}^m \subseteq \mathbf{S}^{NE}$, it follows that, for any two games, such that $\mathbf{S}^{br} = \mathbf{S}^{br'}$, it must be the case that $\mathcal{A}^m = \mathcal{A}^{m'}$. Since ordinal BE solution function picks a subset of $\mathcal{A}^m \subseteq \mathbf{S}^{br}$, it follows that $\phi^o(\mathcal{A}, \mathbf{S}^{br}) = \phi^o(\mathcal{A}', \mathbf{S}^{br'})$ for any two games Γ^o and $\Gamma^{o'}$, such that $\mathbf{S}^{br} = \mathbf{S}^{br'}$.

Since ordinal BE solution is invariant under additions of non-rationalizable outcomes, it also satisfies a version of the *independence of irrelevant strategies axiom*, which requires the solution to be invariant under additions of non-rationalizable strategies.

Every Nash equilibrium of an ordinal game is invariant under order-preserving transformations of ordinal representations of players' preferences. A BE solution is always a Nash equilibrium of the ordinal game, and so it satisfies this *ordinal invariance* axiom.

3.3.4 Examples

In ordinal games with a unique Pareto efficient Nash equilibrium, the BE solution is always the unique Pareto efficient Nash equilibrium. For example, consider the two player coordination game with a unique Pareto efficient Nash equilibrium depicted in Figure 3.5. There are two feasible agreements

	$t1$	$t2$
$s1$	50, 10	70, 20
$s2$	70, 15	60, 5

Figure 3.5: Ordinal coordination game with a unique Pareto efficient Nash equilibrium

in this game – pure strategy Nash equilibria $(s2, t1)$ and $(s1, t2)$. Neither of the two feasible agreements weakly Pareto dominates the other, which means that both agreements are maximally mutually beneficial. Players' losses of attainable individual advantage, as well as differences of players' losses of attainable individual advantage, are shown in Table 3.2. Notice that the row player is indifferent between the two agreements. If any of the two agreements were chosen, the row player would lose no preferred alternatives. The column player, on the other hand, prefers the agreement $(s1, t2)$ over the agreement $(s2, t1)$. If the agreement $(s2, t1)$ were chosen, the column player would forego one preferred alternative. A feasible agreement

Agreement	φ_r^o	φ_c^o	$ \varphi_r^o - \varphi_c^o $
(s1, t2)	0	0	0
(s2, t1)	0	1	1

Table 3.2: Comparison of players' losses of attainable individual advantage

which minimizes the difference between players' losses of attainable individual advantage is the Nash equilibrium (s1, t2). It is the unique ordinal BE solution of this game, which means that hypothetical bargainers would face no coordination problems in this game. The ordinal BE solution function would also resolve coordination problem in other ordinal games with a unique Pareto optimal Nash equilibrium, such as the Hi-Lo game and the Stag Hunt game.

In ordinal games with multiple Pareto efficient pure strategy Nash equilibria, a BE solution may not be unique. Therefore, hypothetical bargainers would face a coordination problem. For example, consider an ordinal version of the Chicken game depicted in Figure 3.6. In this game, each player must

	s	ns
s	40, 4	20, 5
ns	50, 2	10, 0

Figure 3.6: Ordinal Chicken game

simultaneously and independently choose between strategies s (swerve) or ns (not swerve). It has two Nash equilibria in pure strategies – (s, ns) and (ns, s) . Players' losses of attainable individual advantage, as well as differences of players' losses of attainable individual advantage, are shown in Table 3.3: Both Nash equilibria are strictly Pareto optimal, and so both are max-

Agreement	φ_r^o	φ_c^o	$ \varphi_r^o - \varphi_c^o $
(ns, s)	0	2	2
(s, ns)	2	0	2

Table 3.3: Comparison of players' losses of attainable individual advantage

imally mutually advantageous feasible agreements. Notice that both (ns, s) and (s, ns) are the ordinal BE solutions of this game. The game has a BE solution, yet hypothetical bargainers face an action coordination problem. This is an example of a game where ordinal information about players' pref-

ferences is not sufficient to resolve the Nash equilibrium selection problem¹⁵. Other coordination mechanisms, such as, for example, social conventions, are necessary to resolve the coordination problem in this game.

In an ordinal bargaining game with a finite number of possible divisions of resource, a Pareto efficient BE solution will always exist¹⁶. Let us first consider a discrete Divide-the-Cake game with an even number of slices of cake depicted in Figure 3.7. The numbers represent players' ordinal payoffs. In this game, two players are presented with a cake that is cut into four

	<i>0</i>	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>
<i>0</i>	0, 0	0, 1	0, 2	0, 3	0, 4
<i>1</i>	1, 0	1, 1	1, 2	1, 3	0, 0
<i>2</i>	2, 0	2, 1	2, 2	0, 0	0, 0
<i>3</i>	3, 0	3, 1	0, 0	0, 0	0, 0
<i>4</i>	4, 0	0, 0	0, 0	0, 0	0, 0

Figure 3.7: Divide-the-Cake game (even number of pieces)

equal-sized pieces and simultaneously place a demand for the number of pieces for themselves (from 0 to 4). If the sum of their demanded pieces does not exceed 4, they both get what they have asked for. If, on the other hand, the sum exceeds 4, they both get nothing. The game has five Pareto efficient pure strategy Nash equilibria (4, 0), (3, 1), (2, 2), (1, 3), (0, 4) and one *Pareto inefficient* Nash equilibrium (4, 4). Players' losses of attainable individual advantage, as well as the differences of players' losses of attainable individual advantage, are shown in Table 3.4.

A (strictly) Pareto efficient strategy profile (2, 2) is the BE solution of this game. This result usually appeals to most decision-makers, and is supported by experimental results¹⁷.

Let us consider a case where players have to split a larger cake which is cut into five equal-sized pieces. Each of them has to place a demand from 0 to 5 pieces. This discrete Divide-the-Cake game with an odd number of

¹⁵It will later be shown that the cardinal BE solution function can resolve the coordination problem in the cardinal version of the Chicken game.

¹⁶In games with infinite number of possible divisions of resource, the number of feasible agreements is countably infinite, and so the cardinalities of the preferred sets cannot be established with the procedure suggested in this chapter.

¹⁷See Nydegger and Owen (1974) for an experiment in which two players are asked to divide \$1 among themselves and virtually everybody agrees on a 50%-50% split.

Agreement	φ_r^o	φ_c^o	$ \varphi_r^o - \varphi_c^o $
(4, 0)	0	4	4
(3, 1)	1	3	2
(2, 2)	2	2	0
(1, 3)	3	1	2
(0, 4)	4	0	4
(4, 4)	4	4	0

Table 3.4: Comparison of players' losses of attainable individual advantage

pieces is shown in Figure 3.8:

	0	1	2	3	4	5
0	0, 0	0, 1	0, 2	0, 3	0, 4	0, 5
1	1, 0	1, 1	1, 2	1, 3	1, 4	0, 0
2	2, 0	2, 1	2, 2	2, 3	0, 0	0, 0
3	3, 0	3, 1	3, 2	0, 0	0, 0	0, 0
4	4, 0	0, 0	0, 0	0, 0	0, 0	0, 0
5	5, 0	0, 0	0, 0	0, 0	0, 0	0, 0

Figure 3.8: Divide-the-Cake game (odd number of pieces)

This game has six Pareto efficient pure strategy Nash equilibria (5, 0), (4, 1), (3, 2), (2, 3), (1, 4), (0, 5), and one Pareto inefficient Nash equilibrium (5, 5). Players' losses of attainable individual advantage, as well as the differences of players' losses of attainable individual advantage, are shown in Table 3.5:

Agreement	φ_r^o	φ_c^o	$ \varphi_r^o - \varphi_c^o $
(5, 0)	0	5	5
(4, 1)	1	4	3
(3, 2)	2	3	1
(2, 3)	3	2	1
(1, 4)	4	1	3
(0, 5)	5	0	5
(5, 5)	5	5	0

Table 3.5: Comparison of players' losses of attainable individual advantage

The feasible maximally mutually advantageous agreements which minimize the difference between players' losses of attainable individual advantage

are the Nash equilibria $(3, 2)$ and $(2, 3)$. Hypothetical bargainers would identify them both as the BE solutions of this bargaining game. Hypothetical bargaining narrows down the set of solutions, yet hypothetical bargainers would nevertheless face a coordination problem in this game.

3.4 The Cardinal Benefit-Equilibrating Solution

3.4.1 The Intuition Behind the Cardinal BE Solution

To grasp the intuition behind the cardinal BE solution, consider the two player three strategy coordination game depicted in Figure 3.9:

	$t1$	$t2$	$t3$
$s1$	100, 98	0, 0	0, 0
$s2$	0, 0	2, 99	0, 0
$s3$	0, 0	0, 0	1, 100

Figure 3.9: Coordination game with three weakly Pareto optimal outcomes

There are three pure strategy Nash equilibria in this game: $(s1, t1)$, $(s2, t2)$ and $(s3, t3)$. There are also four Nash equilibria in mixed strategies¹⁸. To simplify the presentation of the key principles underlying the cardinal BE solution, in this particular example only the pure strategy Nash equilibria will be considered. Notice that each pure strategy Nash equilibrium is a weakly Pareto optimal outcome. There is a conflict of players' personal interests in this game: The row player's most preferred Nash equilibrium is $(s1, t1)$, which is the least preferred Nash equilibrium for the column player. The column player's most preferred Nash equilibrium is $(s3, t3)$, which is the least preferred Nash equilibrium for the row player. The Nash equilibrium $(s2, t2)$ is, for both players, the second best pure strategy Nash equilibrium in this game.

If the players were hypothetical bargainers and treated this game as an ordinal bargaining problem, they would identify the Nash equilibrium $(s2, t2)$ as the ordinal BE solution of this game: If the Nash equilibrium $(s2, t2)$

¹⁸The four mixed strategy Nash equilibria are: (1) $(\frac{99}{197}s1, \frac{98}{197}s2; \frac{1}{51}t1, \frac{50}{51}t2)$, (2) $(\frac{4950}{14701}s1, \frac{4900}{14701}s2, \frac{4851}{14701}s3; \frac{1}{151}t1, \frac{50}{151}t2, \frac{100}{151}t3)$, (3) $(\frac{50}{99}s1, \frac{49}{99}s3; \frac{1}{101}t1, \frac{100}{101}t3)$, (4) $(\frac{100}{199}s2, \frac{99}{199}s3; \frac{1}{3}t2, \frac{2}{3}t3)$.

were chosen, each player would forego one opportunity to advance personal interests. Given the available information about players' cardinal payoffs, however, such a solution of this game intuitively seems unreasonable: The column player's utility loss seems to be insignificant compared to the utility loss of the row player. In real-world negotiations, the row player could be expected not to accept any agreement other than $(s1, t1)$. If the column player refused, the row player would suffer relatively insignificant payoff losses from, for example, playing his or her mixed *maximin* strategy¹⁹ rather than playing a part in realizing any agreement other than $(s1, t1)$.

Although this intuition is compelling, it is based on an implicit comparison of players' payoffs. In orthodox game theory, players' payoffs are assumed to be von Neumann and Morgenstern utilities. The interpersonal comparisons of von Neumann and Morgenstern utilities are conceptually problematic: The expected utility theory does make numerical representations of individuals' preferences possible, but its principles do not imply their interpersonal comparability (for extensive discussion of why this is so, see, for example, Luce and Raiffa 1957). In other words, the theory offers no answer to the question of how one player's utility units should be 'converted' into utility units of another player.

Although players' utility units may not be interpersonally comparable in the aforementioned sense, there seems to be a conceptual reason to believe that hypothetical bargainers would identify the Nash equilibrium $(s1, t1)$ as the bargaining solution of this game. This would happen if the players were to evaluate the feasible outcomes by comparing their *losses of the maximum attainable individual advantage associated with the implementation of each feasible agreement*.

Assuming that players' cardinal payoffs are common knowledge, each player should be able to identify the personally best outcome for every interacting player. For each player, the personally best outcome is maximally individually advantageous. Any agreement associated with a lower personal payoff than the best outcome can be said to be less individually advantageous than the personally best outcome. If such an agreement were chosen, a player would *suffer a certain loss of maximum attainable individual advantage*. Each player's loss of the maximum attainable individual advantage could be determined if the levels of individual advantage associated with each outcome could be determined.

This, I contend, can be done with a relatively simple Raiffa normalization procedure (Raiffa 1953), also known as the 'zero-one rule', which can be used

¹⁹In this case, the *maximin* strategy of the row player is mixed strategy $(\frac{4950}{14701}s1, \frac{4900}{14701}s2, \frac{4851}{14701}s3)$, yielding the row player a payoff of 100/151.

to measure the level of satisfaction of decision-maker's preferences and which, according to Hausman (1995), is the only legitimate method for comparisons of decision-makers' *preference satisfaction levels* (for extensive discussion, see chapter 2)²⁰. According to this procedure, the level of individual advantage gained from a particular outcome can be defined as the extent by which that outcome advances player's personal payoff from his or her reference point relative to the largest advancement possible, where the latter is associated with the attainment of the outcome that s/he prefers the most.

For the purposes of the BE solution, each hypothetical bargainer's most preferred outcome will be defined as the best possible payoff associated with a rationalizable outcome of the game:

$$u_i^{max} = \max_{s \in S^{br}} u_i(s), \text{ where } S^{br} = (S_1^{br} \times S_2^{br}) \subseteq S. \quad (3.27)$$

As in the case of the ordinal BE solution, two definitions of hypothetical bargainers' reference points seem reasonable. One possibility is to set each hypothetical bargainer's reference point to be the worst personal payoff associated with a *rationalizable outcome* of the game:

$$u_i^{ref} = \min_{s \in S^{br}} u_i(s). \quad (3.28)$$

Another possibility is to set each hypothetical bargainer's reference point to be his or her *maximin payoff level in rationalizable strategies* – a maximum payoff that a player can guarantee to himself or herself, irrespective of which one of the *rationalizable strategies* the opponent is going to choose:

$$u_i^{ref} = \max_{s_i \in S_i} \left\{ \min_{s_{-i} \in S_{-i}^{br}} u_i(s) \right\}. \quad (3.29)$$

For the purposes of the following theoretical discussion, each hypothetical bargainer's reference point will be assumed to be the worst possible payoff associated with a rationalizable outcome of the game.

Consider, again, the coordination game depicted in Figure 3.9. For the row player, the most preferred rationalizable outcome is the Nash equilibrium (s1, t1) (which is also one of the feasible agreements of this game), while the least preferred rationalizable outcome is any outcome of this game associated with a payoff of 0. In line with Raiffa normalization procedure, the value of the row player's most preferred outcome, which will be denoted u_r^{max} , is set to 1, since this outcome leads to the maximum advancement

²⁰Raiffa 1953 and Luce and Raiffa 1957 suggest that the 0-1 normalization is a procedure which can be used to make meaningful interpersonal comparisons of preference satisfaction levels in games where comparisons of players' payoff units are initially not meaningful. See Raiffa 1953 and Luce and Raiffa 1957.

of row player's individual advantage (i.e. the maximum satisfaction of row player's preferences). The value of the least preferred rationalizable outcome, which will be denoted as u_r^{ref} , is set to 0, since it is associated with the lowest advancement of the row player's personal advantage. The levels of individual advantage associated with each of the feasible agreements can be established by applying the appropriate transformation of the row player's original payoffs:

$$u_r^t(\mathbf{s}) = \frac{u_r(\mathbf{s}) - u_r^{ref}}{u_r^{max} - u_r^{ref}}, \text{ where } \mathbf{s} \in \mathbf{S}^{br}. \quad (3.30)$$

It is easy to check that this transformation applied to the row player's reference point will set the value of the reference outcome to 0. The same transformation applied to the row player's most preferred rationalizable outcome (which in this case is the most preferred feasible agreement $(s1, t1)$) will set its value to 1.

The value of the outcome $(s2, t2)$ is a ratio of the utility associated with outcome $(s2, t2)$ to the maximum attainable utility, relative to row player's reference utility:

$$\frac{u_r(s2, t2) - u_r^{ref}}{u_r^{max} - u_r^{ref}} = \frac{2 - 0}{100 - 0} = 0.02. \quad (3.31)$$

The value 0.02 is simply the proportion of the maximum individual advantage that the row player would gain if outcome $(s2, t2)$ were chosen, relative to his or her reference outcome. Since the maximum attainable *level of individual advantage* is 1, the row player would lose, relative to his or her reference outcome, 0.98 of the maximum attainable individual advantage if this outcome were chosen.

The value of the outcome $(s3, t3)$ can be determined with the same procedure:

$$\frac{u_r(s3, t3) - u_r^{ref}}{u_r^{max} - u_r^{ref}} = \frac{1 - 0}{100 - 0} = 0.01. \quad (3.32)$$

The value of the outcome $(s3, t3)$ shows that, relative to reference outcome, the row player would gain 0.01 of the maximum individual advantage attainable in this game. This means that s/he would lose 0.99 of the maximum attainable individual advantage if the outcome $(s3, t3)$ were chosen for implementation.

The same procedure can be used to determine the column player's levels of individual advantage associated with each feasible agreement. For the

column player, the feasible agreement (s_3, t_3) is the best rationalizable outcome. The worst rationalizable outcome is any strategy profile associated with a payoff of 0. By applying the aforementioned procedure, we establish that the values of outcomes (s_1, t_1) , (s_2, t_2) and (s_3, t_3) for the column player are 0.98, 0.99 and 1 respectively. The column player's individual advantage losses associated with outcomes (s_1, t_1) , (s_2, t_2) and (s_3, t_3) are thus 0.02, 0.01 and 0 respectively.

The values determined with the aforementioned procedure can be interpreted in following way. If outcome (s_1, t_1) were chosen, the row player would loose 0 per cent of the maximum attainable individual advantage (i.e. the maximum satisfaction of preferences), while the column player would loose 2 per cent of the maximum attainable individual advantage. If outcome (s_2, t_2) were chosen, the row player would loose 98 per cent of the maximum attainable advantage, while the column player would loose only 1 per cent. Finally, if outcome (s_3, t_3) were chosen, the row player would loose 99 per cent of individual advantage, while the column player would loose none. It is clear that the outcome which minimizes the difference between players' losses of maximum attainable individual advantage is outcome (s_1, t_1) . Outcome (s_1, t_1) is the cardinal BE solution of this coordination game.

Notice that all of the aforementioned measures are established using commonly known and objectively identifiable points in games. When hypothetical bargainers evaluate the feasible agreements, they equate units of measures of their individual advantage – the advancement of their personal interests relative to what they personally deem to be the best and the worst outcome of their interactions – while not being able to equate the units of the attained personal well-being. All they know is how much a particular outcome is individually advantageous to a player relative to that player's reference outcome and his or her most preferred outcome. In order to use these measures, hypothetical bargainers need to know each other's cardinal payoffs and reference points, but they do not need to be able to make interpersonal comparisons of their attained well-being. In other words, the BE solution is a purely *formal arbitration scheme*: It can operate purely on the basis of information about players' reference points and the cardinal payoffs represented by the numbers in the payoff matrix, and so can be used in cases where hypothetical bargainers have no clue as to what kind of personal

motivations those utility numbers actually represent²¹.

3.4.2 Formalization

Let $\Gamma = (\{1, 2\}, \{S_i, u_i\}_{i \in \{1, 2\}})$ be a two player normal form game, in which each player $i \in \{1, 2\}$ has a finite set of pure strategies S_i . Let Σ_i be a set of all probability distributions over the set of pure strategies S_i , and let $\sigma_i \in \Sigma_i$ be a mixed strategy of $i \in \{1, 2\}$, where $\sigma_i(s_i)$ denotes the probability assigned to pure strategy $s_i \in S_i$. Each mixed strategy outcome is a mixed strategy profile $\sigma = (\sigma_1, \sigma_2)$. Each mixed strategy profile should be interpreted as a profile of players' randomized actions. Let $\Sigma = \Sigma_1 \times \Sigma_2$ denote the set of mixed strategy profiles of Γ , and $u_i : \mathcal{L}(\Sigma) \rightarrow \mathbb{R}$ denote the cardinal utility function of player $i \in \{1, 2\}$, which represents i 's preferences over the lotteries over the *set of possible agreements* – the set of mixed strategy profiles of Γ ²².

For each player $i \in \{1, 2\}$, let $\Sigma_i^{br} \subseteq \Sigma_i$ be a set of *rationalizable* strategies, and let $\Sigma^{br} = (\Sigma_1^{br} \times \Sigma_2^{br})$ be a set of *rationalizable strategy profiles* of Γ . Player i 's utility associated with his or her least preferred rationalizable outcome in $\Sigma^{br} \subseteq \Sigma$ will be denoted as $u_i^{ref} := \min_{\sigma \in \Sigma^{br}} u_i(\sigma)$. Player i 's utility associated with his or her most preferred rationalizable outcome in $\Sigma^{br} \subseteq \Sigma$ will be denoted as $u_i^{max} := \max_{\sigma \in \Sigma^{br}} u_i(\sigma)$.

Let $\Sigma^f \subseteq \Sigma^{br}$ be a set of feasible agreements of Γ . As in the ordinal version of the BE solution, it will be assumed that only the self-enforcing agreements are feasible. Since every pure and mixed Nash equilibrium is self-enforcing, the *set of feasible agreements* can be defined as follows:

$$\Sigma^f = \{\sigma \in \Sigma^{br} : \sigma \in \Sigma^{NE}\}, \text{ which implies that } \Sigma^f = \Sigma^{NE}. \quad (3.33)$$

²¹Raiffa normalization is sometimes criticized on the basis of it implicitly ascribing a particular ratio for making interpersonal comparisons of players' payoffs at the point of establishing a common 0–1 scale to legitimize those comparisons. This, it is argued, may lead to inappropriate results when, in some games, the difference between the worst and the best case scenario for one player may be trivial while for another it may be a matter of great importance (see, for example, Sen 1970, Rawls 1971, Griffin 1986 and Hammond 1991). As has been pointed out by Luce and Raiffa 1957 and Hausman 1995, this criticism itself is based on an implicit assumption that players' cardinal utilities represent some objectively comparable levels of welfare, and that comparisons of welfare may not always coincide with comparisons of players' preference satisfaction levels. Yet Raiffa normalization is supposed to be applied to cases where such comparisons are not meaningful in the first place. For extensive discussion, see Luce and Raiffa 1957 and Hausman 1995. This criticism has also been discussed in chapter 2.

²²This definition of players' utility functions represents the assumption that players view the set of mixed strategy profiles as a set of 'goods' over which negotiations take place.

In any game where $u_i^{max} \neq u_i^{ref}$ for every $i \in \{1, 2\}$, the *level of individual advantage* of every player $i \in \{1, 2\}$ associated with any feasible agreement $\sigma \in \Sigma^f$ can be defined as follows:

$$u_i^l(\sigma) = \frac{u_i(\sigma) - u_i^{ref}}{u_i^{max} - u_i^{ref}}. \quad (3.34)$$

Notice that if i 's function is normalized so that $u_i^{max} = 1$ and $u_i^{min} = 0$, then $u_i^l(\sigma) = u_i(\sigma)$.

A feasible agreement $\sigma \in \Sigma^f$ is maximally mutually advantageous *only if* there is no alternative feasible agreement $\sigma' \in \Sigma^f$ which is, for every player $i \in \{1, 2\}$, associated with a strictly higher level of individual advantage than agreement $\sigma \in \Sigma^f$. That is, for every $\sigma \in \Sigma^f$,

$$\sigma \in \Sigma^{fm} \Rightarrow \sigma' \notin \Sigma^f : u_i^l(\sigma') > u_i^l(\sigma) \forall i \in \{1, 2\}^{23}. \quad (3.35)$$

Let $\varphi_i^c(\cdot, \cdot)$ be a *cardinal measure of loss of maximum attainable individual advantage* of player $i \in \{1, 2\}$. A loss of maximum attainable individual advantage associated with a feasible agreement $\sigma \in \Sigma^{fm}$ can, for every player $i \in \{1, 2\}$, be defined in the following way:

$$\varphi_i^c(\sigma, \Sigma^{br}) = \left(\frac{u_i^{max} - u_i^{ref}}{u_i^{max} - u_i^{ref}} \right) - \left(\frac{u_i(\sigma) - u_i^{ref}}{u_i^{max} - u_i^{ref}} \right). \quad (3.36)$$

Since $\frac{u_i^{max} - u_i^{ref}}{u_i^{max} - u_i^{ref}} = 1$ for any $\Sigma^f = \Sigma^{NE}$, the definition of function $\varphi^c(\cdot, \cdot)$ can be simplified in the following way:

$$\varphi_i^c(\sigma, \Sigma^{br}) = 1 - \left(\frac{u_i(\sigma) - u_i^{ref}}{u_i^{max} - u_i^{ref}} \right) = 1 - u_i^l(\sigma). \quad (3.37)$$

The difference between players' losses of maximum attainable individual advantage associated with every $\sigma \in \Sigma^{fm}$ can be defined as follows:

²³As in the ordinal case, this restriction on the set of feasible agreements is based on assumption that self-oriented players would disregard weakly Pareto dominated agreements. The idea is that self-oriented decision-makers will always prefer any weakly Pareto efficient agreement over any weakly Pareto dominated agreement, since every weakly Pareto efficient agreement is associated with a strictly higher utility for every self-oriented decision-maker, irrespective of which weakly Pareto efficient agreement is chosen. A decision to implement a weakly Pareto dominated agreement is incompatible with the assumption that hypothetical bargainers are rational and self-oriented decision-makers – individuals who aim to advance their personal interests as far as possible. For discussion, see Luce and Raiffa 1957, Kalai and Smorodinsky 1975, and Maschler et al. 2013.

$$\left| (1 - u_i^l(\sigma)) - (1 - u_{j \neq i}^l(\sigma)) \right| = \left| \varphi_i^c(\sigma, \Sigma^{br}) - \varphi_{j \neq i}^c(\sigma, \Sigma^{br}) \right|. \quad (3.38)$$

The cardinal BE solution function $\phi^c(\cdot, \cdot)$ satisfies, for every Σ^f ,

$$\phi^c(\Sigma^f, \Sigma^{br}) \in \arg \min_{\sigma \in \Sigma^{fm}} \left\{ \left| \varphi_i^c(\sigma, \Sigma^{br}) - \varphi_{j \neq i}^c(\sigma, \Sigma^{br}) \right| \right\}. \quad (3.39)$$

In cases where some player's preferences over the rationalizable outcomes are such that $u^{max} = u^{ref}$, the proposed 0–1 normalization procedure used to represent the level of player's individual advantage cannot be applied, since the function in such cases is undefined. Such a player is indifferent between all the agreements in $\Sigma^{fm} \subseteq \Sigma^f$, and so there is no agreement in Σ^{fm} which would advance his or her individual interests more than any other agreement in Σ^{fm} . This also implies that an individual cannot view the implementation of any $\sigma \in \Sigma^{fm}$ as being associated with a loss of maximum attainable individual advantage. Therefore, in cases where player's preferences are such that $u^{max} = u^{ref}$, it seems reasonable to assume that $u^l(\sigma) = 1$ for every $\sigma \in \Sigma^{fm}$, and so $\varphi^c(\sigma, \Sigma^{fm}) = 0$ for every $\sigma \in \Sigma^{fm}$.

3.4.3 The Properties of the Cardinal BE Solution

A (possibly non-unique) cardinal BE solution exists in every finite two player non-cooperative game. Nash (1950b, 1951) proved that a Nash equilibrium in mixed strategies exists in every finite game with a finite set of players. In every finite set $\Sigma^f = \Sigma^{NE}$, there must exist at least one $\sigma \in \Sigma^f$, such that

$$\sigma' \notin \Sigma^f : u_i^l(\sigma') > u_i^l(\sigma) \forall i \in \{1, 2\}. \quad (3.40)$$

It follows that $\sigma \in \Sigma^{fm}$, which means that $\Sigma^{fm} \neq \emptyset$, and so a BE solution exists.

Cardinal BE solution is always a *weakly Pareto optimal feasible agreement*. Let $\Sigma^{wpo} \subseteq \Sigma^f$ denote the set of weakly Pareto optimal feasible agreements of Γ . A feasible agreement $\sigma \in \Sigma^f$ belongs to a set of Pareto optimal feasible agreements *only if* there is no alternative feasible agreement $\sigma' \in \Sigma^f$, such that $u_i(\sigma') > u_i(\sigma) \forall i \in \{1, 2\}$. In terms of levels of individual advantage, this condition can be defined as follows:

$$\sigma \in \Sigma^{wpo} \Rightarrow \sigma' \notin \Sigma^f : u_i^l(\sigma') > u_i^l(\sigma) \forall i \in \{1, 2\}. \quad (3.41)$$

Condition (3.41) is equivalent to condition (3.35), which implies that $\Sigma^{wpo} = \Sigma^{fm}$. Since the cardinal BE solution function always picks a subset of $\Sigma^{fm} \subseteq$

Σ^f , it follows that BE solution is always a weakly Pareto optimal feasible agreement.

Cardinal BE solution satisfies a version of the axiom of *individual rationality*, which requires the bargaining solution to be an outcome which, for each player, yields a payoff which is at least as high as his or her *maximin* payoff. The cardinal *maximin* threshold of player $i \in \{1, 2\}$ can be defined as follows:

$$u_i^{mxm} = \max_{\sigma_i \in \Sigma_i} \left\{ \min_{\sigma_{j \neq i} \in \Sigma_j} u_i(\sigma) \right\}. \quad (3.42)$$

The cardinal BE solution satisfies the individual rationality requirement if and only if, for every $i \in \{1, 2\}$,

$$u_i(\sigma) \geq u_i^{mxm} \quad \forall \sigma \in \arg \min_{\sigma \in \Sigma^{fm}} \left\{ \left| \varphi_i^c(\sigma, \Sigma^{br}) - \varphi_{j \neq i}^c(\sigma, \Sigma^{br}) \right| \right\}. \quad (3.43)$$

The *maximin* strategy $\sigma_i^{mxm} \in \Sigma_i$ of player $i \in \{1, 2\}$ is such that

$$u_i(\sigma_i^{mxm}, \sigma_{j \neq i}) \geq u_i^{mxm} \quad \forall \sigma_j \in \Sigma_j. \quad (3.44)$$

If a strategy profile $\sigma \in \Sigma$ is a Nash equilibrium, then the following condition must be satisfied:

$$(\sigma_i, \sigma_{j \neq i}) \in \Sigma^{NE} \Rightarrow u_i(\sigma_i, \sigma_j) \geq u_i(\tilde{\sigma}_i, \sigma_j) \quad \forall \tilde{\sigma}_i \in \Sigma_i. \quad (3.45)$$

It follows that every Nash equilibrium must satisfy the following condition:

$$(\sigma_i, \sigma_{j \neq i}) \in \Sigma^{NE} \Rightarrow u_i(\sigma_i, \sigma_j) \geq u_i^{mxm} \quad \forall i \in \{1, 2\}. \quad (3.46)$$

Since every cardinal BE solution is a Nash equilibrium, the individual rationality requirement is satisfied.

Cardinal BE solution is *invariant under additions of non-rationalizable outcomes*. Suppose that $f^c(\cdot)$ is some cardinal bargaining solution function. The axiom requires that, for any two games Γ and Γ' , such that $\Sigma^{br} = \Sigma^{br'}$, it must be the case that $f^c(\Sigma^{br}) = f^c(\Sigma^{br'})$.

Since u_i^{ref} and u_i^{max} are, for every $i \in \{1, 2\}$, associated with outcomes in $\Sigma^{br} \subseteq \Sigma$, the level of individual advantage $u_i^l(\sigma)$ associated with every $\sigma \in \Sigma^{br}$ remains invariant under additions of outcomes which are such that $\sigma \notin \Sigma^{br}$. For any two games Γ and Γ' , such that $\Sigma^{br} = \Sigma^{br'}$, it must be the case that $\Sigma^{NE} = \Sigma^{NE'}$. Since $\Sigma^{fm} \subseteq \Sigma^{NE}$, it follows that, for any two games, such that $\Sigma^{br} = \Sigma^{br'}$, it must be the case that $\Sigma^{fm} = \Sigma^{fm'}$. Since the cardinal BE solution function picks a subset of $\Sigma^{fm} \subseteq \Sigma^{br}$, it follows that $\phi^c(\Sigma^{fm}, \Sigma^{br}) = \phi^c(\Sigma^{fm'}, \Sigma^{br'})$ for any two games Γ and Γ' , such that $\Sigma^{br} = \Sigma^{br'}$.

Since cardinal BE solution is invariant under additions of non-rationalizable outcomes, it also satisfies a version of the *independence of irrelevant strategies axiom*, which requires the solution to be invariant under additions of non-rationalizable strategies.

The pure strategy Nash equilibria are invariant under positive affine transformations of the interacting players' payoffs: If a pure strategy profile $\mathbf{s} \in \mathbf{S}$ is a Nash equilibrium of Γ , in which player i 's preferences over outcomes are represented by a payoff function $u_i : \mathbf{S} \rightarrow \mathbb{R}$, it will remain a pure strategy Nash equilibrium in any game Γ' , in which i 's preferences over outcomes are represented by a payoff function $u'_i = au_i + c$, where $a > 0$ and c are constants (for a detailed discussion of why this is so, see, for example, Luce and Raiffa, 1957). The mixed strategy Nash equilibria, however, are *not* invariant under positive affine transformations of payoffs. Since a cardinal BE solution of a game may be a mixed strategy Nash equilibrium, this solution concept is not invariant under positive affine transformations of the interacting players' payoffs. It is, however, *invariant under positive scalar transformations of payoffs*. A mixed strategy Nash equilibrium $\sigma \in \Sigma^{NE}$ of a game Γ , in which player i 's preferences are represented by a payoff function $u_i : \mathbf{S} \rightarrow \mathbb{R}$, is a mixed strategy Nash equilibrium in any game Γ' , in which player i 's preferences are represented by a payoff function $u'_i = au_i$, where $a > 0$ (for a technical discussion of why this is so, see, for example, Chen et al. 2009).

3.4.4 Relation to other Bargaining Solutions

The cardinal BE solution shares some conceptual similarities with the Kalai-Smorodinsky bargaining solution (Kalai and Smorodinsky 1975), as well as with the bargaining solution which minimizes the maximum relative concession among the interacting players, which has been suggested by Gauthier (1986).

The Kalai-Smorodinsky bargaining solution can be defined as follows: Suppose that (\mathcal{F}, d) is a two player bargaining problem, where $d = (u_1^{ref}, u_2^{ref})$ is a disagreement point associated with disagreement payoffs of player 1 and player 2 respectively, and \mathcal{F} is the set of feasible agreements. Let u_1^{max} and u_2^{max} denote players' 'ideal payoffs' – the best possible personal payoffs that player 1 and player 2 can attain in the game. The Kalai-Smorodinsky solution function $\mathcal{K}(\mathcal{F}, d)$ picks a point $(u_1, u_2) \in \mathcal{F}$ on a Pareto frontier of \mathcal{F} which maintains the ratio of players' 'ideal' payoff gains:

$$\frac{u_1 - u_1^{ref}}{u_2 - u_2^{ref}} = \frac{u_1^{max} - u_1^{ref}}{u_2^{max} - u_2^{ref}}. \quad (3.47)$$

In the context of non-cooperative games, the Kalai-Smorodinsky bargaining solution could be interpreted as a solution function which picks a feasible agreement on a Pareto frontier of the set of feasible agreements which maintains the ratio of players' 'ideal' payoff gains. That is, the Kalai-Smorodinsky bargaining solution is a strategy profile $\sigma \in \Sigma^{fpf}$, such that

$$\frac{u_1(\sigma) - u_1^{ref}}{u_2(\sigma) - u_2^{ref}} = \frac{u_1^{max} - u_1^{ref}}{u_2^{max} - u_2^{ref}}. \quad (3.48)$$

From definition (3.48) it follows that Kalai-Smorodinsky bargaining solution is a strategy profile $\sigma \in \Sigma^{fpf}$, such that

$$(u_1(\sigma) - u_1^{ref})(u_2^{max} - u_2^{ref}) = (u_2(\sigma) - u_2^{ref})(u_1^{max} - u_1^{ref}). \quad (3.49)$$

A strictly egalitarian cardinal BE solution is a strategy profile $\sigma \in \Sigma^{fm}$, such that

$$\varphi_1^c(\sigma, \Sigma^{br}) = \varphi_2^c(\sigma, \Sigma^{br}). \quad (3.50)$$

Since $\varphi_i^c = 1 - u_i^t(\sigma)$ for every $i \in \{1, 2\}$, it follows that a strictly egalitarian cardinal BE solution is a strategy profile $\sigma \in \Sigma^{fm}$, such that $u_1^t(\sigma) = u_2^t(\sigma)$.

From definition (3.34) it follows that a strictly egalitarian BE solution is a strategy profile $\sigma \in \Sigma^{fm}$, such that

$$\frac{u_1(\sigma) - u_1^{ref}}{u_1^{max} - u_1^{ref}} = \frac{u_2(\sigma) - u_2^{ref}}{u_2^{max} - u_2^{ref}}, \quad (3.51)$$

which is equivalent to

$$(u_1(\sigma) - u_1^{ref})(u_2^{max} - u_2^{ref}) = (u_2(\sigma) - u_2^{ref})(u_1^{max} - u_1^{ref}). \quad (3.52)$$

Notice that (3.52) is equivalent to (3.49). It follows that a strictly cardinal BE solution which is on a Pareto frontier of $\Sigma^{fm} \subseteq \Sigma^{br}$ will have the properties of the Kalai-Smorodinsky bargaining solution. Note, however, that *Kalai-Smorodinsky bargaining solution function cannot be applied to cases where a feasible agreement on a Pareto frontier which maintains the ratios of players' 'ideal' payoff gains does not exist, while the cardinal BE solution function picks a (possibly non-unique) solution in such cases.*

The cardinal BE solution also shares some conceptual similarities with Gauthier's (1986) *minimax relative concession* bargaining solution, which is a point on a Pareto frontier which minimizes the maximum relative concession among the interacting individuals. In the context of non-cooperative games, the *minimax* relative concession solution function could be defined as follows.

Let u_i^{max} and u_i^{ref} denote the ‘ideal’ payoff and the reference payoff of player $i \in \{1, 2\}$ respectively. A *relative concession* of player $i \in \{1, 2\}$ associated with some feasible agreement $\sigma \in \Sigma^f$ can be defined as follows:

$$\frac{u_i^{max} - u_i(\sigma)}{u_i^{max} - u_i^{ref}}. \quad (3.53)$$

A *minimax* relative concession solution is a strategy profile $\sigma \in \Sigma^{fpf}$, such that

$$\sigma \in \arg \min_{\sigma \in \Sigma^{fpf}} \left\{ \max_{i \in \{1, 2\}} \left(\frac{u_i^{max} - u_i(\sigma)}{u_i^{max} - u_i^{ref}} \right) \right\}. \quad (3.54)$$

From definitions (3.34), (3.36) and (3.51), it follows that a strictly egalitarian cardinal BE solution is a strategy profile $\sigma \in \Sigma^{fm}$, such that

$$\left\{ \left(\frac{u_1^{max} - u_1^{ref}}{u_1^{max} - u_1^{ref}} \right) - \left(\frac{u_1(\sigma) - u_1^{ref}}{u_1^{max} - u_1^{ref}} \right) \right\} = \left\{ \left(\frac{u_2^{max} - u_2^{ref}}{u_2^{max} - u_2^{ref}} \right) - \left(\frac{u_2(\sigma) - u_2^{ref}}{u_2^{max} - u_2^{ref}} \right) \right\}, \quad (3.55)$$

which can be simplified to

$$\left(\frac{u_1^{max} - u_1(\sigma)}{u_1^{max} - u_1^{ref}} \right) = \left(\frac{u_2^{max} - u_2(\sigma)}{u_2^{max} - u_2^{ref}} \right). \quad (3.56)$$

From definitions (3.54) and (3.56), it follows that a strictly egalitarian cardinal BE solution minimizes the maximum relative concession among the interacting players. It means that a strictly egalitarian BE solution which is on a Pareto frontier of $\Sigma^{fm} \subseteq \Sigma^{br}$ will have the properties of Gauthier’s *minimax* relative concession bargaining solution. However, *unlike the cardinal BE solution, the minimax relative concession solution does not rely on a measure of the difference of players’ relative concessions*. Therefore, the cardinal BE solution function will often pick a different set of feasible agreements than Gauthier’s *minimax* relative concession solution function. For example, in a game depicted in Figure 3.3(a), Gauthier’s solution function picks outcomes $(s1, t1)$ and $(s3, t3)$, while the cardinal BE solution function picks outcome $(s2, t2)$.

3.4.5 Examples

In some games with multiple weakly Pareto optimal agreements, the cardinal BE solution resolves the coordination problem. For example, consider the following two player three strategy coordination game depicted in Figure

	<i>l</i>	<i>c</i>	<i>r</i>
<i>u</i>	5, 2	0, 0	0, 0
<i>m</i>	0, 0	4, 4	0, 0
<i>d</i>	0, 0	0, 0	2, 5

Figure 3.10: Coordination game with three weakly Pareto optimal outcomes

3.10. This game has three pure strategy Nash equilibria (u, l) , (m, c) and (d, r) . The game also has four Nash equilibria in mixed strategies. Since each of the mixed Nash equilibria is weakly Pareto dominated by at least one pure strategy Nash equilibrium, none of them is in the set of maximally mutually advantageous agreements. For the row player, the maximally individually advantageous outcome is the Nash equilibrium (u, l) , while for the column player the maximally individually advantageous outcome is (d, r) . For both players, the worst outcome is any strategy profile associated with a payoff of 0. The row and the column players' levels of individual advantage, the losses of the maximum attainable individual advantage, as well as the differences between losses of maximum attainable individual advantage associated with each feasible agreement are depicted in Table 3.6:

Agreement	u_r^l	φ_r^c	u_c^l	φ_c^c	$ \varphi_r^c - \varphi_c^c $
(u, l)	1	0	0.4	0.6	0.6
(m, l)	0.8	0.2	0.8	0.2	0
(d, r)	0.4	0.6	1	0	0.6

Table 3.6: Comparison of players' losses of maximum individual advantage

A feasible agreement which uniquely minimizes the difference between players' maximum attainable individual advantage losses is the outcome (m, l) . It is the BE solution of this game. Since this solution is unique, hypothetical bargainers should not face any difficulties coordinating their actions in this game.

In some cases, the BE solution of a game is a mixed strategy Nash equilibrium. Consider the payoff asymmetric Chicken game depicted in Figure 3.11. There are two pure strategy Nash equilibria in this game – (s, ns) and (ns, s) . The third is a mixed strategy Nash equilibrium $(\frac{1}{2}s, \frac{1}{2}ns; \frac{1}{2}s, \frac{1}{2}ns)$, which yields an expected payoff of 3.5 for the row player and 2.5 for the column player. This mixed strategy Nash equilibrium is not weakly Pareto dominated by any other Nash equilibrium, which means that it is in the set of maximally mutually advantageous agreements of this game. For both

	<i>s</i>	<i>ns</i>
<i>s</i>	4, 3	3, 5
<i>ns</i>	6, 2	1, 0

Figure 3.11: Chicken game

players, the worst outcome is (ns, ns) . Notice, however, that it yields different payoffs to the interacting players: The row player's payoff is 1, while the column player's payoff is 0. The row and column player's levels of individual advantage, maximum attainable individual advantage losses, as well as the differences between losses of maximum attainable individual advantage associated with each feasible agreement are depicted in Table 3.7:

Agreement	u_r^i	φ_r^c	u_c^i	φ_c^c	$ \varphi_r^c - \varphi_c^c $
(s, ns)	0.4	0.6	1	0	0.6
(ns, s)	1	0	0.4	0.6	0.6
$(\frac{1}{2}s, \frac{1}{2}ns; \frac{1}{2}s, \frac{1}{2}ns)$	0.5	0.5	0.5	0.5	0

Table 3.7: Comparison of players' losses of maximum individual advantage

The feasible agreement which uniquely minimizes the difference between players' losses of maximum attainable individual advantage is the mixed strategy Nash equilibrium $(\frac{1}{2}s, \frac{1}{2}ns; \frac{1}{2}s, \frac{1}{2}ns)$, and so it is the unique BE solution of this game.

Even relatively simple games may have more than one BE solution. For example, consider the following payoff-symmetric Battle of the Sexes game depicted in Figure 3.12. There are three Nash equilibria in this game: pure

	<i>o</i>	<i>b</i>
<i>o</i>	10, 5	0, 0
<i>b</i>	0, 0	5, 10

Figure 3.12: Battle of the Sexes game

strategy Nash equilibria (o, o) and (b, b) , and a mixed strategy Nash equilibrium $(\frac{2}{3}o, \frac{1}{3}b; \frac{1}{3}o, \frac{2}{3}b)$. The pure strategy Nash equilibria (o, o) and (b, b) are weakly Pareto efficient. The mixed strategy Nash equilibrium yields each player a payoff of $3\frac{1}{3}$. It is therefore weakly Pareto dominated by the pure strategy Nash equilibria of this game. For both players, the worst rationalizable outcome is one of the two outcomes associated with a payoff of 0. The

two pure strategy Nash equilibria are the maximally mutually advantageous agreements in this game. Players' levels of individual advantage, individual advantage losses, as well as the differences between losses of maximum attainable individual advantage associated with each feasible agreement are depicted in Table 3.8:

Agreement	u_r^l	φ_r^c	u_c^l	φ_c^c	$ \varphi_r^c - \varphi_c^c $
(o, o)	1	0	0.5	0.5	0.5
(b, b)	0.5	0.5	1	0	0.5

Table 3.8: Comparison of players' losses of maximum individual advantage

In this payoff symmetric version of the Battle of the Sexes game, both pure strategy Nash equilibria minimize the difference between players' losses of maximum attainable individual advantage. Although a BE solution exists in this game, it is not unique. The hypothetical bargainers would face a coordination problem. It is important to note, however, that the BE solution function resolves the coordination problem in a considerable number of *payoff asymmetric* versions of the Battle of the Sexes game.

3.4.6 Hypothetical Bargaining in Pie Games

One of the fundamental questions pertaining to the hypothetical bargaining theory is whether it can explain real-world decision-maker's behaviour in strategic interactions. A pair of experiments, the design of which seems to be the most suitable for testing the theory, were carried out by Faillo et al. (2013, 2016). In both experiments, the participants played two player coordination games, known as 'pie games', in which they had to choose one of the three outcomes represented as segments of a pie. An example of a pie game is shown in Figure 3.13. If we call the top left slice $r1$, the top right slice $r2$, and the bottom slice $r3$, then the pairs $(r1, r1)$, $(r2, r2)$ and $(r3, r3)$ can be viewed as outcomes yielding pairs of payoffs $(9, 10)$, $(10, 9)$ and $(9, 9)$ to the players respectively. The labels $r1$, $r2$ and $r3$ were hidden from participants and the positions of the slices were varied across three different treatment groups (see Figure 3.14). If both participants chose the same outcome (i.e. managed to coordinate their actions), they received positive payoffs. A normal form representation of this game is provided in Figure 3.15. The type of pie games used and the conclusions drawn in the two experiments were fairly similar.

The original experiments were designed to pit the team reasoning the-




Figure 3.13: An example of a pie game from Faillo et al. (2013, 2016) experiments (Source: Faillo et al. 2013)




Figure 3.14: An example of a pie game depicted in Figure 3.13 as seen by the two interacting players in one treatment. The positions of the three slices were varied across treatments (Source: Faillo et al. 2013)

ory²⁴ against the cognitive hierarchy theory²⁵. However, the structure of pie

²⁴For an extensive discussion of the team reasoning theory, see chapters 1 and 2.

²⁵The cognitive hierarchy theory postulates the existence of players of different cognitive levels. Each cognitive level represents the degree to which that type of player can reason about the other players. The level 0 decision-makers choose any of the available strategies at random (i.e. they play each available strategy with equal probability). The level 1 reasoners assume everybody else to be level 0 players and best respond to the level 0 player's strategy. The level 2 reasoners assume everybody else to be either level 1 or level 0 player (in simpler versions of the theory, level 2 player believes that everyone else is a level 1 player) and, similarly, best respond to the expected strategies of level 1 and/or level 0 players (in simpler models, always to level 1 player's strategy). The same logic applies to players of higher cognitive levels. Although in principle the cognitive hierarchy theory allows for any number of cognitive types (where each type assumes the other players to be of a lower cognitive level than themselves), in practice it is usually assumed that most decision-makers are level 1 or level 2 reasoners. For details, see Nagel 1995, Stahl and Wilson 1995, Ho et al. 1998, Camerer et al. 2004, Crawford et al. 2008, Bardsley et al. 2010, Faillo et al. 2016.

	<i>r1</i>	<i>r2</i>	<i>r3</i>
<i>r1</i>	9, 10	0, 0	0, 0
<i>r2</i>	0, 0	10, 9	0, 0
<i>r3</i>	0, 0	0, 0	9, 9

Figure 3.15: 3x3 pie game depicted in Figure 22 represented in normal form

games is particularly suitable to test the theory of hypothetical bargaining, and even compare it with the team reasoning function based on the notion of mutual advantage suggested in chapter 2. Notice that pie games share substantial structural similarities with standard bargaining games, such as the Divide-the-Cake game. The players have to choose between several different allocations of payoffs and, in case their choices disagree, they do not receive a positive payoff. In addition, each allocation of payoffs is a Nash equilibrium. Since the labels are hidden and the positions of slices can be varied, the choices of players are less likely to be influenced by factors other than the payoff structure of the game itself (i.e. the number of the available allocations of payoffs and their properties). Table 3.9 summarizes the results of Faillo et al. (2013). All the outcomes which satisfy the properties of a cardinal BE solution (i.e. would be selected by the cardinal BE solution function) are indicated by be . The outcomes which maximize the mutual advantage of team-reasoning decision makers (i.e. would be selected by the team reasoning function based on the notion of mutual advantage suggested in chapter 2) are indicated by $_{\tau}$. Assuming that hypothetical bargainer/team

	G1	G2	G3	G4	G5	G6	G7	G8	G9	G10	G11
<i>r1</i>	9, 10	9, 10	9, 10	9, 10	10, 10	10, 10	10, 10	10, 10	9, 12	10, 10	9, 11
<i>r2</i>	10, 9	10, 9	10, 9	10, 9	10, 10	10, 10	10, 10	10, 10	12, 9	10, 10	11, 9
<i>r3</i>	9, 9	11, 11	9, 8	11, 10	9, 9	11, 11	9, 8	11, 10	10, 11	11, 9	10, 10
$N(\%)r1$	14_{τ}	0	51_{τ}^{be}	16	48_{τ}^{be}	1	51_{τ}^{be}	26_{τ}	16	43_{τ}^{be}	6
$N(\%)r2$	11_{τ}	1	45_{τ}^{be}	4	34_{τ}^{be}	3	31_{τ}^{be}	22_{τ}	11	27_{τ}^{be}	7
$N(\%)r3$	74_{τ}^{be}	99_{τ}^{be}	4	80_{τ}^{be}	18	96_{τ}^{be}	18	52_{τ}^{be}	73_{τ}^{be}	30	86_{τ}^{be}

Table 3.9: Summary of Faillo et al 2013 results. Choices predicted by the cardinal BE model are indicated by be . Choices predicted by the team reasoning model are indicated by $_{\tau}$

reasoning decision-maker always chooses to play his or her part in the attainment of a mutually advantageous outcome (in case of multiple solutions, s/he chooses to play a part in the attainment of one of the multiple mutually advantageous solutions at random), we can compare the theoretical ‘predic-

	G1	G2	G3	G4	G5
l	5, 6	5, 10	5, 6	5, 6	6, 7
r	6, 5	10, 5	6, 5	6, 5	7, 6
b	5, 5	5, 5	6, 5	7, 5	7, 5
$N(\%)l(P1); N(\%)l(P2)$	6; 7 $_{\tau}$	0; 13 $_{\tau}$	53; 21 $_{\tau}^{be}$	40; 38 $_{\tau}$	35; 33 $_{\tau}^{be}$
$N(\%)r(P1); N(\%)r(P2)$	6; 0 $_{\tau}$	7; 13 $_{\tau}$	16; 33 $_{\tau}^{be}$	35; 29 $_{\tau}^{be}$	40; 33 $_{\tau}^{be}$
$N(\%)b(P1); N(\%)b(P2)$	88; 93 $_{\tau}^{be}$	93; 73 $_{\tau}^{be}$	32; 46 $_{\tau}^{be}$	25; 33 $_{\tau}$	0, 14

Table 3.10: Summary of Crawford et al 2010 results. The choices of player 1 (P1) and player 2 (P2) are presented separately. Choices predicted by the cardinal BE model are indicated by be . Choices predicted by the team reasoning model are indicated by $_{\tau}$

tions’ of the hypothetical bargaining model with the theoretical ‘predictions’ of the model of team reasoning based on the notion of mutual advantage.

In Faillo et al. 2013 experiment, the aforementioned model of team reasoning based on the notion of mutual advantage is a reasonably good predictor of choices in 9 out of 11 games (G2, G3, G4, G5, G6, G7, G9, G10, G11). The cardinal BE solution is a good predictor of choices in 10 out of 11 games (G1, G2, G3, G4, G5, G6, G7, G8, G9, G11).

In a series of experiments carried out by Crawford et al. (2008), the participants were presented with five pie games. The players had to choose one of the three weakly Pareto optimal outcomes: (l, l) , (r, r) and (b, b) . It seems reasonable to assume that the asymmetry of payoffs present in some of the pie games made the comparison of outcomes a more cognitively demanding task.

Table 3.10 summarizes the results of Crawford et al. (2008). The choices of Player 1 and Player 2 are presented separately. All the outcomes which satisfy the properties of a cardinal BE solution are indicated by be . The outcomes which maximize the mutual advantage of team-reasoning decision makers are indicated by $_{\tau}$.

The results suggest that a hypothetical bargaining model based on the cardinal BE solution concept is a reasonably good predictor of participants’ choices in 4 out of 5 games – G1, G2, G3 and G5. In game G4, outcome (r, r) is the unique cardinal BE solution, and so the theory fails to predict a considerable number of participants choosing outcomes (l, l) and (b, b) ²⁶. The theory of team reasoning based on the notion of mutual advantage suggested in chapter 2 offers reasonably good predictions in 3 out of 5 games:

²⁶It can be argued that this model offers extremely weak theoretical predictions in game G3, since every outcome is the BE solution.

G3, G4 and G5²⁷.

Given the assumption that hypothetical bargainers/team-reasoning players always choose to play a part in the attainment of a mutually beneficial solution (in games with multiple solutions, one of the multiple mutually beneficial solutions at random), the hypothetical bargaining model based on the cardinal BE solution seems to offer slightly more accurate theoretical predictions of decision-makers' choices in the aforementioned games than the team reasoning theory based on the notion of mutual advantage. It is important to note, however, that both the hypothetical bargaining theory and the theory of team reasoning based on the notion of mutual advantage are theories of how players *identify* the mutually advantageous solutions of games, not how such decision-makers *actually choose their strategies*. As such, neither of the two theories offers an answer to the question of how and under what conditions the players should be expected to play their part in the attainment of a mutually advantageous outcome. More importantly, both theories do not offer a theoretically sound answer to the question of how players coordinate their actions in games with multiple mutually advantageous solutions. For either theory to be adequately tested, the model must be complemented with an *epistemic model of choice*, specifying the conditions under which a player should be expected to play his or her part in the attainment of a mutually advantageous outcome, as well as with a model of how players resolve coordination problems in games with multiple mutually advantageous solutions. However, since the coordination aids used by the interacting decision-makers may be based on shared cultural norms and social practices, a general formal game theoretic model of such players' *final strategy choices* in games may not be possible. For an extensive discussion of the aforementioned problems, see chapters 2 and 4.

3.5 Possible Extensions of the Model

3.5.1 Application to N -Player Games

So far, the discussion focused on the question of how the principles of the bargaining theory could be applied in the game theoretic analysis of non-cooperative two player games. The BE solution concept suggested in this chapter can be applied to any two player finite game. However, another important question pertaining to the hypothetical bargaining theory is whether

²⁷It can be argued that this model offers extremely weak theoretical predictions in games G3 and G4, since every outcome maximizes the mutual advantage of the interacting individuals.

it can be used to analyze games with more than two players.

It is clear that the model can be easily applied to any game which has a unique Pareto optimal feasible agreement, since hypothetical bargainers will identify such an agreement as the unique BE solution of the game. In games with multiple weakly Pareto optimal feasible agreements, however, the task of identifying the most equitable distribution of individual advantage losses (in ordinal case, the foregone preferred alternatives) becomes much more complex.

In games with a strictly egalitarian feasible maximally mutually advantageous agreement, the identification of the BE solution is unproblematic. In most games, however, a strictly egalitarian BE solution will not exist. Nevertheless, the hypothetical bargainers could still distinguish the maximally individually advantageous (i.e. weakly Pareto optimal) feasible agreement associated with a *more equitable* distribution of individual advantage losses (foregone preferred alternatives in the ordinal case) from the one associated with a less equitable distribution of individual advantage losses by making *pairwise comparisons* of players' losses of maximum individual advantage.

Recall that, in games with two players, the BE solution is identified by comparisons of feasible agreements based on the differences between the row player's and the column player's losses of the maximum attainable individual advantage. In games with more than two players, the difference between the losses of maximum attainable individual advantage can be determined for any possible pair of players.

Suppose that a set $I = (1, \dots, m)$ of hypothetical bargainers are playing a normal form finite game with a set $\Sigma^{fm} \subseteq \Sigma^{NE}$ of maximally mutually advantageous agreements. It is clear that, for any two players $i \in I$ and $j_{j \neq i} \in I$, the difference between their losses of the maximum attainable individual advantage associated with the implementation of each maximally mutually advantageous feasible agreement $\sigma \in \Sigma^{fm}$ can be determined in the same way as in the two player case:

$$\left| (1 - u_i^t(\sigma)) - (1 - u_{j_{j \neq i}}^t(\sigma)) \right|. \quad (3.57)$$

Any two feasible agreements $\sigma \in \Sigma^{fm}$ and $\sigma' \in \Sigma^{fm}$ should be deemed indistinguishable in terms of comparisons of losses of maximum attainable individual advantage if, for every pair of players, the difference between their losses of maximum attainable individual advantage associated with each agreement is the same. An agreement $\sigma \in \Sigma^{fm}$ should be deemed more egalitarian than any agreement $\sigma' \in \Sigma^{fm}$ if, *for every pair of players*, agreement $\sigma \in \Sigma^{fm}$ is associated with a strictly smaller difference between their losses of maximum attainable individual advantage than agreement

$\sigma' \in \Sigma^{fm}$. That is, an agreement $\sigma' \in \Sigma^{fm}$ should be *ruled out* as a possible cardinal BE solution if, *for every possible pair of players* $i \in I$ and $j_{j \neq i} \in I$,

$$\left| (1 - u'_i(\sigma)) - (1 - u'_j(\sigma)) \right| < \left| (1 - u'_i(\sigma')) - (1 - u'_j(\sigma')) \right| \quad (3.58)$$

For example, consider a three player game depicted in Figure 3.16. In this

	<i>t1</i>	<i>t2</i>		<i>t1</i>	<i>t2</i>
<i>s1</i>	10, 9, 9	0, 0, 0	<i>s1</i>	5, 5, 5	0, 0
<i>s2</i>	0, 0, 0	5, 5, 5	<i>s2</i>	0, 0	6, 8, 10
	<i>m1</i>			<i>m2</i>	

Figure 3.16: Three player coordination game played by three hypothetical bargainers

game, the row player chooses between strategies *s1* and *s2*, the column player chooses between strategies *t1* and *t2*, and the matrix player chooses between strategies *m1* and *m2*. To simplify the example, it will be assumed that players only consider pure strategy outcomes.

This game has two weakly Pareto optimal pure strategy Nash equilibria: $(s1, t1, m1)$ and $(s2, t2, m2)$. The hypothetical bargainers should therefore identify both outcomes as the maximally mutually advantageous feasible agreements of this game. For each player, the worst rationalizable outcome is any strategy profile associated with a payoff of 0. Players' levels of individual advantage, their losses of maximum attainable individual advantage associated with each feasible agreement, and the differences between players' losses of maximum attainable individual advantage are shown in Table 3.11:

Agreement	u'_r	φ'_r	u'_c	φ'_c	u'_m	φ'_m	$ \varphi'_r - \varphi'_c $	$ \varphi'_r - \varphi'_m $	$ \varphi'_c - \varphi'_m $
$(s1, t1, m1)$	1	0	1	0	$\frac{9}{10}$	$\frac{1}{10}$	0	$\frac{1}{10}$	$\frac{1}{10}$
$(s2, t2, m2)$	$\frac{6}{10}$	$\frac{4}{10}$	$\frac{8}{9}$	$\frac{1}{9}$	1	0	$\frac{13}{45}$	$\frac{4}{10}$	0

Table 3.11: Players' levels of individual advantage, losses of maximum individual advantage, and differences of individual advantage losses

Table 3.11 shows the differences between the row player's and the column player's, the row player's and the matrix player's, and the column player's and the matrix player's losses of the maximum attainable individual advantage. Notice that, for each pair, the outcome $(s1, t1, m1)$ is associated with a strictly lower difference of losses of maximum attainable individual

advantage than the outcome $(s2, t2, m2)$. Therefore, the Nash equilibrium $(s1, t1, m1)$ would be identified as the one associated with the most equitable distribution of losses of maximum attainable individual advantage, and so it is the cardinal BE solution of this game. Notice that this result is intuitively compelling. If the outcome $(s2, t2, m2)$ were chosen, the matrix player would lose no maximum attainable individual advantage, while the row player and the column player would lose 40% and 20% of their maximum attainable individual advantage attainable in this game respectively. If, on the other hand, the outcome $(s1, t1, m1)$ were chosen, the row player would lose no maximum attainable individual advantage, while the row and the column player would each lose 10% of the maximum attainable individual advantage. The differences between players' losses of maximum attainable individual advantage associated with the outcome $(s2, t2, m2)$ are clearly larger than those associated with the outcome $(s1, t1, m1)$.

There is another way in which the more equitable agreements can be distinguished from the less equitable ones. Notice that a strictly egalitarian cardinal BE solution is a feasible agreement, such that players' losses of the maximum attainable individual advantage are equal. This implies that players' individual advantage levels associated with a strictly egalitarian BE solution must also be equal. That is, a strictly egalitarian BE solution is a strategy profile $\sigma \in \Sigma^{fm}$, such that $u_i^t(\sigma) = u_{j \neq i}^t(\sigma)$, for every pair $i, j \in I$, where $I = (1, \dots, m)$ is the set of players. If we were to sum up the individual advantage levels of every interacting player in some m -player game, the ratio of each player's level of individual advantage to the sum of players' individual advantage levels should always be $1/m$.

Let $\sum_{i \in I} u_i^t(\sigma)$ be the sum of individual advantage levels associated with a strategy profile $\sigma \in \Sigma^{fm}$ of every $i \in I$. If a strategy profile $\sigma \in \Sigma^{fm}$ is a strictly egalitarian BE solution, it must necessarily be the case that

$$\frac{u_i^t(\sigma)}{\sum_{i \in I} u_i^t(\sigma)} = \frac{1}{m}, \text{ for every } i \in I. \quad (3.59)$$

In many games, a strictly egalitarian cardinal BE solution will not exist. However, the equity of any two feasible maximally mutually advantageous agreements can be evaluated by comparing each player's *actual* ratio of individual advantage level to the sum of players' individual advantage levels associated with each agreement with an ideal strictly egalitarian ratio of $\frac{1}{m}$. That is, for any maximally mutually advantageous feasible agreement $\sigma \in \Sigma^{fm}$ and every player $i \in I$, we can determine the distance between the *actual* ratio of player's individual advantage level to the sum of all individuals' advantage levels and the *ideal egalitarian ratio*:

$$\left| \frac{u_i^t(\sigma)}{\sum_{i \in I} u_i^t(\sigma)} - \frac{1}{m} \right| \geq 0. \quad (3.60)$$

Any two agreements $\sigma \in \Sigma^{fm}$ and $\sigma' \in \Sigma^{fm}$ can then be compared by comparing the distance between the actual ratio and the ideal egalitarian ratio of $1/m$ of every player $i \in I$. Any agreement $\sigma \in \Sigma^{fm}$ should be deemed more egalitarian than any agreement $\sigma' \in \Sigma^{fm}$ if, for every $i \in I$,

$$\left| \frac{u_i^t(\sigma)}{\sum_{i \in I} u_i^t(\sigma)} - \frac{1}{m} \right| < \left| \frac{u_i^t(\sigma')}{\sum_{i \in I} u_i^t(\sigma')} - \frac{1}{m} \right|^{28}. \quad (3.61)$$

For example, consider, again, the game depicted in Figure 3.16. Let $U_{\{r,c,m\}}^t = u_r^t(\mathbf{s}) + u_c^t(\mathbf{s}) + u_m^t(\mathbf{s})$ denote the sum of players' individual advantage levels associated with some $\mathbf{s} \in \mathbf{S}^{fm}$. If this game had an egalitarian cardinal BE solution, it would be some $\mathbf{s} = (s_r, s_c, s_m)$, such that

$$\frac{u_r^t(\mathbf{s})}{U_{\{r,c,m\}}^t(\mathbf{s})} = \frac{u_c^t(\mathbf{s})}{U_{\{r,c,m\}}^t(\mathbf{s})} = \frac{u_m^t(\mathbf{s})}{U_{\{r,c,m\}}^t(\mathbf{s})} = \frac{1}{3}. \quad (3.62)$$

The sum of players' individual advantage levels associated with the agreement $(s1, t1, m1)$ is 2.9. The sum of players' individual advantage levels associated with the agreement $(s2, t2, m2)$ is 2.48889. The ratio of each player's individual advantage level to the sum of players' individual advantage levels, and the distance between each player's actual ratio and the hypothetical strictly egalitarian cardinal BE solution ratio are shown in Table 3.12. The numbers are rounded off to 3 decimal places.

Agreement	$\frac{u_r^t}{U_{\{r,c,m\}}^t}$	$\frac{u_c^t}{U_{\{r,c,m\}}^t}$	$\frac{u_m^t}{U_{\{r,c,m\}}^t}$	$\left \frac{u_r^t}{U_{\{r,c,m\}}^t} - \frac{1}{3} \right $	$\left \frac{u_c^t}{U_{\{r,c,m\}}^t} - \frac{1}{3} \right $	$\left \frac{u_m^t}{U_{\{r,c,m\}}^t} - \frac{1}{3} \right $
$(s1, t1, m1)$	0.345	0.345	0.310	0.011	0.011	0.023
$(s2, t2, m2)$	0.241	0.357	0.402	0.092	0.023	0.069

Table 3.12: The ratios of each players individual advantage level to the sum of players individual advantage levels, and the distance between each players actual ratio and the strictly egalitarian ratio

It is easy to see that, for each player, the outcome $(s1, t1, m1)$ is associated with a smaller difference between the ratio of his or her individual advantage level to the sum of players' individual advantage levels and the

²⁸In any n -player ordinal game, a feasible agreement $\mathbf{s} \in \mathcal{A}^m$ is more egalitarian than $\mathbf{s}' \in \mathcal{A}^m$ if, for every $i \in \{1, \dots, m\}$,

$$\left| \frac{C_i(\mathbf{s}, \mathcal{A})}{\sum_{i \in \{1, \dots, m\}} C_i(\mathbf{s}, \mathcal{A})} - \frac{1}{n} \right| < \left| \frac{C_i(\mathbf{s}', \mathcal{A})}{\sum_{i \in \{1, \dots, m\}} C_i(\mathbf{s}', \mathcal{A})} - \frac{1}{n} \right|.$$

ideal strictly egalitarian ratio of 1/3 than the outcome $(s2, t2, m2)$. Since the outcome $(s1, t1, m1)$ is associated with a more equitable distribution of individual advantage gains than the outcome $(s2, t2, m2)$, hypothetical bargainers should identify the agreement $(s1, t1, m1)$ as the cardinal BE solution of this game.

The preceding discussion relied on the assumption that every player of the game is a hypothetical bargainer, and that this fact is common knowledge among the interacting players. In a two player game, it seems reasonable to assume that a hypothetical bargainer will not search for the BE solution of the game if s/he does not believe that the opponent will do that as well. More complicated problems arise in m -player games when some of the players are hypothetical bargainers while others are not. If hypothetical bargainers were aware of the fact that some of the players are not hypothetical bargainers, or were uncertain about the players' type, they could still consider playing their part in realizing the BE solution of the game by taking into account their predictions of the strategy choices of players who are not reasoning as hypothetical bargainers. However, the strategy choices of the players who are not reasoning as hypothetical bargainers may render the implementation of the BE solution impossible.

For example, consider the three player coordination game depicted in Figure 3.17. This game has two weakly Pareto optimal pure strategy Nash

	$t1$	$t2$		$t1$	$t2$
$s1$	10, 9, 10	0, 0, 0	$s1$	4, 4, 5	0, 0
$s2$	0, 0, 0	5, 5, 5	$s2$	0, 0	8, 10, 10
	$(m1)$			$(m2)$	

Figure 3.17: Three player coordination game, in which the matrix player is not a hypothetical bargainer

equilibria: $(s1, t1, m1)$ and $(s2, t2, m2)$. For each player, the worst rationalizable outcome is one of the pure strategy outcomes associated with a payoff of 0. For hypothetical bargainers, the cardinal BE solution of this game is the Nash equilibrium $(s1, t1, m1)$.

Suppose that it is common knowledge among the row and the column player that they are hypothetical bargainers, but they have no information about the matrix player's type. They could not attain the outcome $(s1, t1, m1)$ without the matrix player choosing strategy $m1$.

If the row and the column player were to believe that the matrix player will choose $m1$, they would have a reason to believe that their choices of

strategies s_1 and t_1 will actually lead to the attainment of the outcome (s_1, t_1, m_1) . Notice, however, that the matrix player is indifferent between the Nash equilibria (s_1, t_1, m_1) and (s_2, t_2, m_2) . The payoff structure of the game does not provide any cues as to which strategy the matrix player is more likely to choose.

In such a situation of strategic uncertainty, the row and the column player could consider playing a combination of strategies (s_2, t_2) , since it guarantees each player a minimum payoff of 5, irrespective of what the matrix player does. Notice that if the players were to play a combination of strategies (s_1, t_1) , they could only guarantee themselves a minimum payoff of 4.

This example shows that the BE solution may not be chosen in a strategic situation, in which some of the players are not hypothetical bargainers, or in a strategic situation, in which hypothetical bargainers are uncertain about each other's type. It also shows that hypothetical bargainers must hold fairly complex shared beliefs in order to be motivated to implement the BE solution²⁹.

3.5.2 Social Interactions and the Set of Feasible Agreements

The BE solution proposed in this chapter is based on the assumption that the set of agreements deemed feasible by hypothetical bargainers is the set of the Nash equilibria of a game. From the perspective of orthodox game theory, this assumption seems reasonable: If a hypothetical bargainer believes that his or her opponents are individualistically rational decision-makers, s/he should expect them to deviate from any agreement which creates them a personal incentive to do so. Therefore, a hypothetical bargainer should not expect his or her opponent to play a part in realizing a strategy profile if that opponent's agreement strategy is not a best response to the agreement strategies of other hypothetical bargainers. If rationality is common knowledge, hypothetical bargainers should expect each other to deviate from non-self-enforcing agreements, and so only deem an agreement feasible if it creates no 'double crossing' incentives (for an extensive discussion of the role of feasibility criterion in bargaining theory, see, for example, Myerson 1991).

In real-world social interactions, however, social agents may be using a less restrictive feasibility criterion. For example, a decision-maker may only be concerned about the personal payoff losses, and so deem an agreement feasible if opponents' deviations cannot lead to a loss of his or her personal

²⁹The structure of beliefs which is required to motivate the hypothetical bargainers to implement the BE solution will be discussed in considerable detail in chapter 4.

payoff. If decision-makers were to use such a feasibility criterion, some of the non-equilibrium strategy profiles would be identified as feasible agreements.

Misyak and Chater (2014) have suggested this possibility, and, to test this hypothesis, conducted an experiment with the Boobytrap game depicted in Figure 3.18. Like the Prisoner's Dilemma game, the Boobytrap game has

	<i>cooperate</i>	<i>defect</i>	<i>boobytrap</i>
<i>cooperate</i>	30, 30	10, 40	30, 29
<i>defect</i>	40, 10	20, 20	-100, 9
<i>boobytrap</i>	29, 30	9, -100	29, 29

Figure 3.18: The Boobytrap Game

a unique pure strategy Nash equilibrium (*defect, defect*). However, each player has an additional option of playing the *boobytrap* strategy. If one of the players chooses this strategy, the other player's best response is strategy *cooperate*, while the choice of the strategy *defect* yields the player the worst personal payoff in this game. If both players choose the strategy *boobytrap*, each of them gets a payoff higher than the one associated with the Nash equilibrium (*defect, defect*).

Strategy profile (*boobytrap, boobytrap*) is not a Nash equilibrium: Both players have an incentive to deviate by choosing the strategy *cooperate*. A unilateral or even bilateral individually advantageous deviations from the strategy profile (*boobytrap, boobytrap*) do not harm any of the players: If one of the players were to choose to deviate by playing strategy *cooperate*, the player who were to play the strategy *boobytrap* would get the same payoff of 29, while the deviating player would get a slightly better payoff of 30. If both players were to choose to deviate by playing strategy *cooperate*, both of them would get a payoff of 30.

The problem with the temptation to deviate by playing *cooperate* is that an opponent expecting such a deviation will have an incentive to deviate himself or herself by playing *defect*. According to Misyak and Chater (2014), this could be prevented if both players were to play a mixed strategy of 'cooperative choices' *boobytrap* and *cooperate*, in which strategy *boobytrap* is played with a sufficiently high probability (in this particular game with a probability infinitesimally higher than 1/14) to deter the opponent from playing strategy *defect*. In that case, the players could prevent each other from playing strategy *defect*, and would attain a more mutually beneficial outcome than by playing the Nash equilibrium (*defect, defect*): Under the former solution, the expected payoff of each player would very nearly be 30

– much better than the payoff associated with the latter. The results from Misyak and Chater’s experiments with the boobytrap game seem to suggest that combinations of cooperation and boobytrap strategies constitute 82% of game outcomes (for an extensive discussion of experimental results, see Misyak and Chater 2014).

From the perspective of orthodox game theory, however, Misyak and Chater’s suggestion is conceptually problematic. To see why, notice that even the aforementioned profile of players’ mixed strategies is an *inherently unstable agreement*: Each player will be motivated to deviate from the mixed strategy profile by not playing the strategy *boobytrap* at all, since randomized strategy where the pure strategy *boobytrap* is played with a positive probability yields him or her a strictly lower expected payoff than the pure strategy *cooperate*. In other words, each self-oriented (i.e. expected utility maximizing) player will be motivated not to implement a strategy that eliminates the opponent’s incentive to choose strategy *defect*. An opponent expecting such a deviation will be motivated to play strategy *defect*. If players were self-oriented, and this fact were common knowledge, they would expect each other to deviate from the agreement, and respond optimally by playing strategy *defect*. An agreement to play the aforementioned randomized strategy profile would thus unravel.

Although problematic from the perspective of orthodox game theory, Misyak and Chater’s suggestion cannot be ruled out as a *description of what strategy profiles the real-world decision-makers view as feasible agreements*. If Misyak and Chater’s suggestion is correct, then in certain games the set of feasible agreements considered by real-world hypothetical bargainers may be larger than the set of the Nash equilibria. Hypothetical bargaining could thus potentially explain how people identify mutually advantageous agreements in certain games with inefficient Nash equilibria. For this reason, Misyak and Chater’s suggestion warrants further empirical testing.

3.6 Conclusion

Misyak and Chater’s virtual bargaining theory offers a conceptually innovative explanation of how self-oriented decision-makers resolve games with multiple rationalizable outcomes. However, I have argued that theory’s reliance on the standard Nash bargaining solution for modelling hypothetical bargaining in non-cooperative games is conceptually problematic. My arguments focused on the standard Nash bargaining solution function’s insensitivity to alternative feasible distributions of players’ personal payoff gains.

In this chapter I have suggested a model of hypothetical bargaining based

on the benefit-equilibrating bargaining solution concept, the principles of which are broadly in line with the principles underlying the ordinal egalitarian bargaining solution concept for finite sets of Pareto optimal points suggested by Conley and Wilkie (2012). I have argued that the hypothetical bargaining model based on the suggested BE solution function offers a plausible theoretical explanation of how players resolve conflicts over alternative allocations of individual gains in non-cooperative games, in which the interpersonal comparisons of their payoffs are assumed not to be meaningful.

In my account I primarily focused on the question of what properties an outcome must have in order to be identified by hypothetical bargainers as the BE solution of a game. The question of how hypothetical bargainers coordinate their actions in games with multiple BE solutions warrants further investigation and discussion. For example, consider the extended Hi-Lo game depicted in Figure 3.19(a). There are two BE solutions of this game – pure strategy Nash equilibria $(hi1, hi1)$ and $(hi2, hi2)$. Hypothetical bargainers would face a coordination problem. The model proposed in this chapter does not offer an answer to the question of how hypothetical bargainers should coordinate their actions in this game: In terms of the formal properties, both solutions are equivalent.

	<i>hi1</i>	<i>hi2</i>	<i>lo</i>		<i>hi1</i>	<i>hi2</i>	<i>lo</i>
<i>hi1</i>	10, 10	0, 0	0, 0	<i>Hi1</i>	10, 10	0, 0	0, 0
<i>hi2</i>	0, 0	10, 10	0, 0	<i>Hi2</i>	0, 0	10, 10	0, 0
<i>lo</i>	0, 0	0, 0	9, 9	<i>Lo</i>	0, 0	0, 0	5, 5
	(a)				(b)		

Figure 3.19: Two versions of the extended Hi-Lo game

It must be stressed that the suggested bargaining model should not be viewed as a coordination theory, but rather as a theory of how players may use the commonly known information about the payoff structure of the game in identifying a feasible and mutually beneficial solution. In games where the BE solution is unique, its identification resolves the coordination problem for hypothetical bargainers. In games with multiple BE solutions, however, hypothetical bargainers could take multiple approaches towards resolving the coordination problem. For example, they could choose their BE solution strategies randomly (that is, they could play each of their BE strategies with equal probabilities). Alternatively they could take into consideration the *perceived coordination success rate*, and then consider playing *ex ante*

weakly Pareto dominated feasible agreements if, given the coordination success rate, the *ex ante* weakly Pareto dominated feasible agreement yields a higher expected payoff for every bargainer than the *ex ante* weakly Pareto optimal feasible agreement³⁰. For example, in the aforementioned extended Hi-Lo game depicted in Figure 3.19(a), hypothetical bargainers could choose the outcome (lo, lo) which, although *ex ante* not maximally mutually advantageous, is unique. By playing their part in realizing outcome (lo, lo) , both players could gain a guaranteed payoff of 9, which is higher than the expected payoff of 5 associated with the aforementioned randomized choices of strategies *hi1* and *hi2*.

The model suggested in this chapter could, in principle, be modified to include coordination success rates into the formal characterization of the BE solution. However, this *ad hoc* modification of the solution concept is conceptually problematic. First, the solution could not be taken to represent the outcome of the mental simulation of the actual bargaining process, since in open negotiations the bargainers would definitely agree on playing either the Nash equilibrium $(hi1, hi1)$ or the Nash equilibrium $(hi2, hi2)$.

Second, such a modification of the solution concept would not resolve the coordination problem in every possible scenario. For example, consider a version of the extended Hi-Lo game depicted in Figure 3.19(b).

In this game, the players would get the same expected payoff from randomizing between strategies *hi1* and *hi2* as they would get from playing the Nash equilibrium (lo, lo) . The question of what the players should choose to do in such situations of strategic uncertainty cannot be answered with the tools of the theory suggested in this chapter, and further research into the psychological factors that influence players' belief formation process may be necessary to explain coordination in such games. Since players' ability to coordinate their actions may often depend on factors that are not related to payoff structures of games alone, a single generalizable formal model of their *final* choices may not be possible³¹.

³⁰In the context of team reasoning theory, this idea has been discussed by Bardsley et al. 2010 and Faillo et al. 2016.

³¹For an in-depth discussion of this problem, see chapters 2 and 4.

Chapter 4

Hypothetical Bargaining, Social Coordination and Evolution

4.1 Introduction

In most general terms, orthodox game theory can be defined as the study of mathematical models of conflict and cooperation between two or more perfectly rational decision-makers. The proponents of a position known as ‘methodological dualism’ defend the view that a clear distinction can and ought to be maintained between normative and descriptive game theory (see, for example Aumann 1985). The normative game theory explores the ‘nature and the consequences of idealized full rationality in strategic interactions’, while the descriptive game theory ‘aims at the explanation and prediction of observed behavior’. (Selten 1988: vii). The question of empirical adequacy is of fundamental importance for descriptive game theory, but less so for the normative game theory.

As has been pointed out by Selten, the distinction between the normative and the descriptive game theory, is ‘blurred in the practice of applied research’, since the ‘methods developed in normative theory are used in the analysis of applied models in the hope for empirical relevance’ (Selten 1988: vii). In other words, the analytic methods and solution concepts used in the normative analysis are also used in the formal models pertaining to explain and/or predict the behaviour of social agents.

In recent decades, descriptive game theory has become an important analytic tool of social sciences: Various more or less complex interactions between two or more social agents are modelled as games played by rational agents. The theoretical predictions of perfectly rational decision-makers’ actions in the idealized models of strategic interactions are taken to be infor-

mative of the actions of social agents in real-world interdependent decision problems.

Despite its widespread use, the orthodox game theoretic analysis has certain conceptual limitations, one of them being the indeterminacy of its theoretical predictions in non-cooperative games with multiple Nash equilibria (Bacharach and Bernasconi 1997). Even the simplest of games which are interpreted by social scientists as idealized models of real-world social interactions have multiple Nash equilibria. From the perspective of orthodox game theory, every Nash equilibrium of a game is a rational solution, and so games with multiple Nash equilibria have multiple rational solutions. Yet some of the games with multiple Nash equilibria seem to have intuitively ‘obvious’ *unique* rational solutions. Experimental evidence suggests that such games create virtually no coordination problems for real-world decision-makers. The orthodox theory of games, however, offers no compelling explanation of why people identify one Nash equilibrium as an ‘obvious’ solution of a game rather than any other Nash equilibrium, or why people expect each other to play the ‘obvious’ solution. An example of a game where the standard game theoretic analysis leads to conclusions which contradict our common sense intuitions about rationality is the three strategy common interest game depicted in Figure 4.1:

	$r1$	$r2$	$r3$
$r1$	9, 10	0, 0	0, 0
$r2$	0, 0	10, 9	0, 0
$r3$	0, 0	0, 0	11, 10

Figure 4.1: Common interest game

There are three pure strategy Nash equilibria in this game: Strategy profiles $(r1, r1)$, $(r2, r2)$ and $(r3, r3)$. In addition, the game has four Nash equilibria in mixed strategies¹. From the perspective of orthodox game theory, all the Nash equilibria are rational solutions of this game. Yet for many people the Nash equilibrium $(r3, r3)$ stands out as a more compelling solution of this game than any other Nash equilibrium. It is a strategy profile which, for each player, is associated with the highest personal payoff attainable in this game. Intuitively it seems that intelligent players who understand

¹The four mixed strategy Nash equilibria of this game are: $(\frac{9}{19}r1, \frac{10}{19}r2; \frac{10}{19}r1, \frac{9}{19}r2)$, $(\frac{9}{28}r1, \frac{5}{14}r2; \frac{9}{28}r3; \frac{110}{299}r1, \frac{99}{299}r2, \frac{90}{299}r3)$, $(\frac{1}{2}r1, \frac{1}{2}r3; \frac{11}{20}r1, \frac{9}{20}r3)$, $(\frac{10}{19}r2, \frac{9}{19}r3; \frac{11}{21}r2, \frac{10}{21}r3)$.

the structure of this game should expect each other to play the Nash equilibrium $(r3, r3)$ and choose their strategies which are consistent with this expectation (for extensive discussion, see Bacharach 2006). Results from experiments support this intuition by revealing that approximately 80% of people opt for strategy $r3$ (for experimental results, see, for example, Faillo et al. 2013).

The critics have argued that the failure of the standard game theory to agree with our ‘high quality’ intuitions about the solutions of common interest games, similar to the one depicted in Figure 4.1, indicates its failure to capture all the relevant strategic considerations of real-world decision-makers, such as collective optimality considerations, and so should be viewed as a conceptual limitation of descriptive game theory (see for example, Olcina and Urbano 1994, Bacharach 2006). Some of the critics have suggested that a development of a model of reasoning which better approximates the actual process of reasoning by which people arrive at their strategy choices is the appropriate response to this problem (see, for example, Bacharach 2006, Misyak and Chater 2014).

One of the more recent theories which offers an answer to this question is the hypothetical, or virtual, bargaining theory (Misyak and Chater 2014, Misyak et al. 2014), which suggests that players choose their actions in non-cooperative games on the basis of what they believe they would agree to play if they could openly bargain – engage in real negotiations, in which each player can communicate his or her offers to the other players and receive their counteroffers. When a player reasons as a hypothetical, or virtual, bargainer, s/he views the pure and mixed strategy profiles, or outcomes, as possible agreements that players could implement via joint actions. S/he then identifies a set of *feasible agreements* – a subset of agreements, such that each agreement in that subset is a pure or mixed strategy profile which does not allow any of the players to exploit the other player(s) by unilaterally deviating from it. The player then identifies a feasible agreement which, s/he believes, the players would agree on playing in open bargaining, and plays his or her part in realizing it, provided that s/he has a reason to believe that the other players are hypothetical bargainers and will carry out their end of the agreement. The strategy profile that hypothetical bargainers believe they would agree to play in open negotiations is the *mutually beneficial and agreeable* resolution of the game (for details, see Misyak and Chater 2014, Misyak et al. 2014).

Hypothetical bargaining theory shares many conceptual similarities with Sugden’s recent version of the theory of team reasoning based on the notion

of mutual advantage (Sugden 2011, 2015)². Team reasoning is a theory that certain structural and/or contextual features of games may trigger a shift in decision-maker's mode of reasoning from individualistic best-response reasoning to reasoning as a member of a team – a group of individuals who act together in the attainment of some common goal. When a person reasons as a member of a team, s/he identifies a strategy profile – one strategy for each player in a game – that leads to the attainment of the best possible outcome for the group of individuals acting together as a team. The team-reasoning decision-maker then chooses a strategy which, in combination with other team-reasoning decision-makers' strategies, leads to the attainment of that outcome.

In Sugden's (2011, 2015) recent version of the team reasoning theory, an outcome selected by decision-makers who reason as members of a team must be perceived as being mutually advantageous by every interacting team-reasoning decision-maker. Sugden proposes to define mutually advantageous outcomes as those that are associated with decision-makers' personal payoffs satisfying a particular threshold. The suggested threshold is each player's personal *maximin* payoff – the maximum payoff that a player can guarantee to himself or herself in a particular game irrespective of what the other players are going to do (for details, see Sugden 2015).

Both theories can be viewed as the so-called *goal-directed reasoning theories*: A decision-maker accepts certain premises about the decision problem and his or her goal, and then follows a set of well-defined inference rules in order to identify an action that s/he should take in order to make the attainment of the goal possible³. Both theories are also mutual advantage-directed reasoning theories, since they both suggest that decision-makers aim to resolve games by identifying outcomes which they perceive as being mutually advantageous. In other words, both theories suggest that the goal of decision-makers in non-cooperative games is the identification and attainment of mutually advantageous outcomes – strategy profiles which, relative to each decision-maker's reference point, advance his or her personal interests. In fact, some of the proponents of the hypothetical bargaining theory suggest that hypothetical bargaining can be viewed as a complement of the team reasoning theory. According to Misyak and Chater, hypothetical, or virtual, bargaining theory 'can be viewed as providing a link between indi-

²For early developments of this theory see Sugden (1993, 2000, 2003) and Bacharach (1999, 2006). For some of the more recent work see Gold and Sugden (2007a,b), Sugden (2011, 2015) and Gold (2012).

³For an extensive discussion of goal-directed reasoning models, see, for example, Gold and List 2004, Bacharach 2006, Gold and Sugden 2007b, Smerilli 2014, Misyak and Chater 2014, Misyak et al. 2014.

vidual beliefs and values and the behaviour of “the team” – by viewing the preferences of the team as resulting from what “the team” would agree, if its members had the opportunity to bargain.’ (Misyak and Chater 2014: 4). This interpretation of hypothetical bargaining, however, may not be acceptable to all of the proponents of team reasoning, since team reasoning is sometimes interpreted as a mode of reasoning which only characterizes the reasoning of players who identify themselves with a team due to overlapping personal interests (see, for example, Bacharach 2006, Zizzo and Tan 2007 and Smerilli 2014), while hypothetical bargaining can occur between ‘adversaries’ – self-oriented individuals with incompatible private interests (Misyak et al. 2014).

Hypothetical bargaining theory has been introduced primarily as a social coordination account – a theory which purports to explain how social agents coordinate their actions in coordination problems where they cannot communicate and cannot use a commonly known social rule of behaviour to coordinate their actions. The proponents of the theory argue that hypothetical bargaining is the origin of various ‘unwritten rules’ of social interactions: Hypothetical bargainers ‘compose’ the rules for social interactions as they go along, by engaging in a process of mental bargaining. In other words, the proponents of this theory suggest that people do not need to know the pre-existing rule of social behaviour (e.g. a social norm or a convention) in order to be able coordinate their actions in an interdependent decision problem: They can figure out the appropriate response by identifying the hypothetical bargaining solution of the game. That is, when a hypothetical bargainer identifies an outcome as a bargaining solution of the game, s/he also identifies a strategy that s/he must choose in order to make the attainment of that outcome possible. If every player of a game is a hypothetical bargainer and chooses the appropriate strategy, hypothetical bargainers end up playing the outcome which they identify as the bargaining solution of the game. Thus, hypothetical bargainers’ ability to identify the same outcome of the game as the hypothetical bargaining solution resolves the coordination problem (for extensive discussion, see Misyak et al. 2014).

On the surface of it, it might seem that the theory of hypothetical bargaining is a parsimonious social coordination theory. Unlike, for example, the social convention theory (see, for example, Bicchieri 2006 and Gintis 2008), which presupposes that players know the appropriate behavioural rule for every type of interdependent decision problem, expect each other to know it, expect each other to follow it, and so on, the hypothetical bargaining theory offers an explanation of how people can coordinate their actions in a large variety of games with a relatively simple reasoning algorithm – a reasoning

procedure which allows each decision-maker who uses it to determine a solution of a game from a commonly known information about decision-makers' preferences over the possible outcomes. It could also be viewed as providing a credible explanation of why one individual's deviation from the expected pattern of behaviour trigger a negative response from other individuals: If decision-makers view an implementation of a specific strategy profile as an implicit 'agreement' leading to the attainment of a mutually beneficial resolution of a game, then they may also view each deviation from that strategy profile as a 'violation' of their hypothetical agreement – a deviation from a 'pre-agreed' pattern of behaviour, which warrants a punitive response.

In this chapter I will argue that although the hypothetical bargaining theory offers a relatively parsimonious explanation of how people identify the payoff salient solutions in a large variety of non-cooperative interdependent decision problems, at best it offers only a partial explanation of how people coordinate their actions in non-cooperative games. I will focus on two epistemic limitations of the theory. In section 2 I will discuss the epistemic assumptions of the hypothetical bargaining model, and argue that, since the theory of hypothetical bargaining is supposed to be interpreted as a model of rational decision-making, it is vulnerable to the *rationalization problem*: The model of hypothetical bargaining does not fully account for the structure of beliefs which sustains hypothetical bargainer's motivation to play his or her part in the implementation of a hypothetical agreement. I will discuss several responses to this problem and point out their limitations. In section 3 I will argue that hypothetical bargaining, if interpreted as a rational social coordination theory, is vulnerable to the problem of common beliefs: The theory cannot account for the structure of beliefs which makes it a functioning social coordination mechanism. In section 4 I will discuss a possible evolutionary game theoretic response to this problem, and highlight its explanatory limitations. In section 5 I will argue that even a fully developed hypothetical bargaining theory would not provide a single generalizable model of players' final choices due to non-uniqueness of hypothetical bargaining solutions. With section 6 I conclude.

4.2 Hypothetical Bargaining and Rational Choice

4.2.1 Hypothetical Bargaining as a Reasoning Algorithm

The proponents of the theory suggest that hypothetical bargaining 'operates within the framework of rational-choice theory' (Misyak et al. 2014: 512). In other words, they seem to suggest that hypothetical bargaining is a model

of strategic decision-making which is compatible with the principles of the standard game and decision theory.

In standard game theory models, a rational player is assumed to be a best-response reasoner – a decision-maker who always chooses a strategy which, given his or her beliefs about the opponents' strategy choices, maximizes expected payoff. In a hypothetical bargaining model, the players are assumed to be hypothetical bargainers – decision-makers who choose their strategies on the basis of what they would agree to play if they could openly bargain. In other words, a hypothetical bargainer views non-cooperative games as bargaining problems, and aims to resolve them by identifying a strategy profile, or outcome, which *s/he believes* the players would agree to play in real negotiations, in which players can communicate their offers and counteroffers to each other until an agreement is reached (Misyak and Chater 2014). An outcome identified by hypothetical bargainers as the hypothetical bargaining solution of a game is the expected outcome of an explicit bargaining process. According to the proponents of the theory, hypothetical bargaining process mimics the explicit bargaining process with no external enforcement of agreements, and so idealized game theoretic models of explicit bargaining can be used as idealized models of hypothetical bargaining. In other words, the principles and solution concepts used in the analysis of game theoretic models of explicit bargaining can be used in the game theoretic models of hypothetical bargaining (Misyak and Chater 2014, Misyak et al. 2014).

Given the aforementioned assumptions, hypothetical bargaining could be interpreted as the following reasoning algorithm. Each hypothetical bargainer views the set of mixed and pure outcomes of a game as the set of possible agreements that bargainers could implement by engaging in joint actions. Hypothetical bargainer then identifies a set of feasible agreements. There are several suggestions of how the notion of feasibility could be defined as a theoretical concept. In standard bargaining models without external enforcement, the set of feasible agreements is assumed to be the set of *correlated equilibria* of the game (for extensive discussion of why this is so, see Myerson 1991 and Maschler et al. 2014). In non-cooperative games, however, decision-makers are assumed not to be able to communicate at all, which means that an agreement to implement a correlated equilibrium which requires a correlation device is impossible. Therefore, the set of feasible agreements of a non-cooperative game could be defined as the set of its

Nash equilibria⁴.

Misyak and Chater (2014) suggest an alternative and less restrictive definition of feasibility, according to which a strategy profile is a feasible agreement if no player can exploit the other players by varying his or her strategy and gaining individual advantage to the disadvantage of others. If hypothetical bargainers were to adopt this feasibility criterion, then in certain games the set of agreements which they would identify as feasible would be larger than the set of Nash equilibria (for discussion, see chapter 3). In the following discussion, the set of feasible agreements will be assumed to be the set of Nash equilibria. This assumption will not make the following discussion any less general, since the problems of the hypothetical bargaining model based on assumption that the set of feasible agreements is the set of Nash equilibria will also be present in any hypothetical bargaining model based on a less restrictive feasibility criterion.

After determining the set of feasible agreements, a hypothetical bargainer identifies a feasible agreement which s/he believes the players would most likely agree to play in open bargaining. By identifying the hypothetical bargaining solution of a game, each hypothetical bargainer also identifies a strategy that s/he must play in order to make the implementation of that solution possible. If hypothetical bargainer expects the other bargainers to play their part in the implementation of the identified bargaining solution, s/he will be motivated to play his or her part in the implementation of that solution as well.

One of the central questions pertaining to hypothetical bargaining theory is what properties a feasible agreement must have in order to be identified by hypothetical bargainers as the bargaining solution of a non-cooperative game. Misyak and Chater (2014) suggest the Nash bargaining solution (Nash 1950a) as a reasonably good approximation to what hypothetical bargainers would identify as the hypothetical bargaining solution of a game, since this solution ‘follows from very simple and natural axioms concerning bargaining’ (Misyak and Chater 2014: 4)⁵.

As I have argued in chapter 3, the Nash bargaining solution may not be the best solution concept to represent hypothetical bargaining in non-cooperative games, since it is insensitive to information about the possible alternative allocations of players’ personal payoff gains. Because of these rea-

⁴Every Nash equilibrium is also a correlated equilibrium, which means that the set of Nash equilibria in a non-cooperative game is a subset of the set of its correlated equilibria. For extensive discussion, see Aumann 1987 and Brandenburger and Dekel 1987.

⁵The Nash bargaining solution satisfies the following axioms: Pareto optimality, invariance with respect to affine utility transformations, independence of irrelevant alternatives and symmetry. For extensive discussion, see Nash 1950a.

sons, other bargaining solution concepts, such as the benefit-equilibrating bargaining solution suggested in chapter 3, may serve as better approximations to what hypothetical bargainers would identify as the bargaining solution of a game. However, the following discussion will focus on the general epistemic assumptions of the hypothetical bargaining theory, and so the arguments will apply to any model of hypothetical bargaining, irrespective of which bargaining solution concept is chosen as an approximation to what hypothetical bargainers identify as the bargaining solution of a game. Therefore, the games used as examples in the following discussion will be such that the outcome which will be assumed to be the hypothetical bargaining solution of a game will be in line with the Nash bargaining solution (Nash 1950a), the Kalai-Smorodinsky bargaining (Kalai and Smorodinsky 1975) solution⁶, as well as with the benefit-equilibrating bargaining solution suggested in chapter 3.

4.2.2 The Rationalization Problem

Every game theoretic model of strategic interactions relies on certain assumptions about players' beliefs. The hypothetical bargaining models are no exception. A simple model of hypothetical bargaining, such as the one suggested by Misyak and Chater (2014), relies on assumption that every decision-maker is a hypothetical bargainer and this fact is common knowl-

⁶The Kalai-Smorodinsky bargaining solution can be defined as follows: Suppose that (\mathcal{F}, d) is a two player bargaining problem, where $d = (u_1^{ref}, u_2^{ref})$ is a disagreement point associated with disagreement payoffs of player 1 and player 2 respectively, and \mathcal{F} is the set of feasible agreements. Let u_1^{max} and u_2^{max} denote players' 'ideal payoffs' – the best possible personal payoffs that player 1 and player 2 can attain in the game. The Kalai-Smorodinsky solution function $\mathcal{K}(\mathcal{F}, d)$ picks a point $(u_1, u_2) \in \mathcal{F}$ on a Pareto frontier of \mathcal{F} which maintains the ratio of players' 'ideal' payoff gains:

$$\frac{u_1 - u_1^{ref}}{u_2 - u_2^{ref}} = \frac{u_1^{max} - u_1^{ref}}{u_2^{max} - u_2^{ref}}.$$

The Kalai-Smorodinsky bargaining solution satisfies the following axioms: Pareto optimality, symmetry, invariance with respect to affine utility transformations and monotonicity. Unlike the Nash bargaining solution (Nash 1950a), the Kalai-Smorodinsky bargaining solution does not satisfy the independence of irrelevant alternatives axiom. For extensive discussion, see Kalai and Smorodinsky 1975.

edge⁷ among the interacting individuals. The model also relies on a relative standard assumption that the payoff structure of the game is common knowledge. In other words, it is assumed that each hypothetical bargainer knows that every other player is a hypothetical bargainer who knows the payoff structure of the game, knows that every other player knows that every other player is a hypothetical bargainer who knows the payoff structure of the game, and so on *ad infinitum*. Assuming that decision-makers only deem possible one type of hypothetical bargainer (i.e. the players believe that every hypothetical bargainer is using the same reasoning algorithm), the aforementioned common knowledge assumptions imply that each player knows that every player of the game has identified the same outcome as the bargaining solution of the game, knows that every other player knows it, and so on *ad infinitum*⁸. Every player also knows that every player has identified a strategy that s/he must choose in order to play a part in the attainment of that outcome, knows that every other player knows this, and so on *ad infinitum*. If a hypothetical bargainer expects the other bargainers to carry out their part in the attainment of the identified outcome, s/he will choose to play his or her part in the attainment of that outcome as well. In other words, a hypothetical bargainer will play a part in the attainment of the outcome only if s/he expects the others to do the same.

The fundamental problem of this model is revealed by the following observation: Hypothetical bargainer's knowledge of the fact that every player has identified a certain outcome as the hypothetical bargaining solution does not give him or her a valid (i.e. rational) reason to believe that the other hypothetical bargainers will play their part in the attainment of that outcome, even if the identified outcome is a Nash equilibrium. In other words, even if every hypothetical bargainer identifies a certain Nash equilibrium as the hypothetical bargaining solution of a game, and this fact is common

⁷In more recent epistemic models of game theory, a *common belief* in rationality assumption tends to be preferred over the common knowledge of rationality assumption due to its purported realism. Common knowledge of rationality assumption is viewed as a stronger epistemic assumption than common belief in rationality assumption: Common knowledge assumption implies that the player cannot be mistaken in his or her belief, while a common probability 1 belief may be either true or false. For the purposes of brevity, the standard common knowledge term will be used in most cases, since the conceptual differences between the common knowledge and the common belief concepts will not play any significant role in the following discussion (that is, the terms can be used interchangeably without undermining the validity of arguments). For an extensive discussion of the differences between common knowledge and common belief concepts, see Samet 2013.

⁸The coordination problems which arise in games with multiple hypothetical bargaining solutions will be discussed separately in section 4.5.

knowledge, each player’s decision to play his or her part in the attainment of that Nash equilibrium cannot be explained as an expected utility maximizing choice by appeals to the fact that everyone has identified that equilibrium as the bargaining solution and that this fact is common knowledge. It is only rational for a player to play his or her part in the Nash equilibrium if s/he expects the opponent to do that as well. Yet if the decision-maker believes that the opponent is rational, s/he should expect the opponent to play his or her part in the Nash equilibrium only if s/he has a reason to believe that the opponent expects the decision-maker to play his or her part in that equilibrium himself or herself. Hypothetical bargainers who believe in each other’s rationality thus end up in an infinite regress which gives neither of them a valid reason for playing his or her part in the attainment of the Nash equilibrium which they identify as the bargaining solution⁹. Thus, common knowledge of the fact that a certain outcome is the hypothetical bargaining solution of a game is *not sufficient to rationalize hypothetical bargainer’s decision to play a part in the attainment of that outcome*. A similar problem would obviously arise with any out-of-equilibrium outcome which satisfies the feasibility criterion suggested by Misyak and Chater (2014).

If hypothetical bargaining model is to be interpreted as a model explaining players’ *actual strategy choices* rather than merely as a *model of reasoning*, its epistemic assumptions must be strengthened considerably. More specifically, common knowledge of the fact that every player is a hypothetical bargainer must be complemented with a common p -belief that each bargainer will play his or her part in the attainment of the outcome identified as the bargaining solution. In other words, hypothetical bargainers must express a common belief that each hypothetical bargainer will choose to play a part in the attainment of that outcome with probability of at least $p \geq 0$, which is high enough to make the choice of playing a part in the attainment of that outcome optimal for every interacting hypothetical bargainer¹⁰.

There are two possible ways in which the model could be modified to satisfy the aforementioned epistemic requirement. The simplest modification of the theory would be to treat the hypothetical bargaining as a ‘mechanistic’ choice algorithm rather than as a belief-based choice algorithm. That is, a

⁹This argument is similar to Gilbert’s (1989) criticism of Schelling’s (1960) theory of salience. Gilbert’s argument rests on a simple observation that player’s decision to play a salient equilibrium cannot be justified as being rational by appeals to the fact that the Nash equilibrium is recognized as salient by every player and this recognition is common knowledge. The players must have a reason to expect each other to actually choose the strategies constituting the salient equilibrium. For details, see Gilbert 1989.

¹⁰For a technical characterization of the common p -belief concept, see Monderer and Samet 1989 and Kajii and Morris 1997.

hypothetical bargainer could be assumed to be a decision-maker who always plays his or her part in the attainment of an outcome that s/he identifies as the bargaining solution of a game (i.e. players commitment to the attainment of that outcome is unconditional). Such a mechanistic model of choice would, contrary to claims made by Misyak et al. 2014, obviously involve a non-trivial departure from the standard game theoretic model of strategic reasoning, which rests on the principle that each player's strategy choice depends on his or her preferences over outcomes, as well as probabilistic beliefs about the opponents' strategy choices. Although such a model of hypothetical bargaining is relatively unproblematic as a *descriptive* model of choice, no claims about its *normative* appeal could be justified.

Another, and more subtle, approach is to treat the hypothetical bargaining reasoning procedure as a *belief-formation algorithm*: A reasoning procedure by which the decision-maker forms a belief about the combination of strategies that the other players are going (or most likely going) to play. In other words, a strategy profile that a hypothetical bargainer identifies as the hypothetical bargaining solution is simply a hypothetical bargainer's belief about the strategy profile that the other hypothetical bargainers will (or most likely will) be implementing. If a hypothetical bargainer expects everyone else to play their part in a particular Nash equilibrium, then playing his or her part in that equilibrium is an optimal response.

For example, consider a normal form game $\Gamma = (I \{S_i, u_i\}_{i \in I})$, where $I = \{1, \dots, m\}$ is the set of players, S_i is the set of pure strategies and $u_i : \mathbf{S} \rightarrow \mathbb{R}$ is the payoff function of player $i \in I$, where $\mathbf{S} = \times_{i \in I} S_i$ is the set of pure strategy profiles of Γ . Each pure strategy profile $\mathbf{s} = (s_1, \dots, s_m)$ is a combination of player's pure strategies – one for each player of the game. For every game Γ , it is possible to define a set $S_{-i} = S_1 \times \dots \times S_{i-1} \times S_{i+1} \times \dots \times S_m$ which represents all the possible combinations of strategies of all the players other than player $i \in I$. Each combination of strategies $s_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_m)$ is a combination of strategies of all the players other than player i . Each pure strategy profile $\mathbf{s} \in \mathbf{S}$ can then be written as a tuple $\mathbf{s} = (s_i, s_{-i})$, where $s_{-i} \in S_{-i}$ is a combination of components (pure strategies) of all the players other than i and $s_i \in S_i$ is i 's component of the strategy profile $\mathbf{s} \in \mathbf{S}$.

Let Σ_i be a set of probability distributions over S_i . Each mixed strategy $\sigma_i \in \Sigma_i$ is a particular probability distribution over player i 's pure strategies, where $\sigma_i(s_i)$ denotes the probability assigned to pure strategy $s_i \in S_i$. Let $\Sigma = \times_{i \in I} \Sigma_i$ be the set of all the mixed strategy profiles of Γ , and $u_i(\sigma) = \sum_{\mathbf{s} \in \mathbf{S}} (\prod_{i \in I} \sigma_i(s_i)) u_i(\mathbf{s})$ be player i 's expected utility associated with a mixed strategy profile $\sigma \in \Sigma$, where $\sigma = (\sigma_1, \dots, \sigma_m)$ is a particular combination of players' mixed strategies. Similarly as in the pure strategy case, we can

define a set $\Sigma_{-i} = \Sigma_1 \times \dots \times \Sigma_{i-1} \times \Sigma_{i+1} \times \dots \times \Sigma_m$ of all the possible combinations of mixed strategies of every player other than player $i \in I$, where each element $\sigma_{-i} \in \Sigma_{-i}$ is a particular combination of mixed strategies of all the players other than player $i \in I$. A mixed strategy profile $\sigma \in \Sigma$ can be written as a tuple $\sigma = (\sigma_i, \sigma_{-i})$, where $\sigma_{-i} \in \Sigma_{-i}$ is a combination of mixed strategies of all the players other than i and $\sigma_i \in \Sigma_i$ is i 's component of the strategy profile $\sigma \in \Sigma$.

The hypothetical bargaining could be interpreted as a belief-formation algorithm in the following way. Suppose that every player $i \in I$ is a hypothetical bargainer, and this is common knowledge among the interacting players. Suppose that hypothetical bargainer identifies a strategy profile $\sigma^* \in \Sigma$ as the *unique* hypothetical bargaining solution of Γ . By identifying outcome $\sigma^* \in \Sigma$ as the hypothetical bargaining solution, hypothetical bargainer $i \in I$ forms a probability 1 belief that every hypothetical bargainer will play his or her part in the combination of strategies $\sigma_{-i}^* = (\sigma_1^*, \dots, \sigma_{i-1}^*, \sigma_{i+1}^*, \dots, \sigma_m^*)$. Since the fact that every player is a hypothetical bargainer is common knowledge, each player $i \in I$ believes that the probability of every other player of that game playing his or her part in the combination of strategies σ_{-i}^* is 1. Given this structure of beliefs, hypothetical bargainer $i \in I$ will be motivated to play his or her part $\sigma_i^* \in \Sigma_i$ in the profile $\sigma^* \in \Sigma$ if and only if

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*) \quad \forall \sigma_i \in \Sigma_i. \quad (4.1)$$

Notice that if $\sigma^* \in \Sigma$ is a Nash equilibrium, then i 's strategy $\sigma_i^* \in \Sigma_i$ is a best response to $\sigma_{-i}^* \in \Sigma_{-i}$. Assuming that hypothetical bargainer always forms a probability 1 belief that the other hypothetical bargainiers will play their part in the attainment of an outcome identified as the bargaining solution of a game, and that the set of agreements which hypothetical bargainiers deem feasible is always the set of the Nash equilibria of a game, a hypothetical bargainer will always play a part in the implementation of the hypothetical bargaining solution, provided that it is common knowledge among the interacting players that every player of the game is a hypothetical bargainer.

The suggested interpretation of hypothetical bargaining can be shown to share certain similarities with some of the more recent epistemic models of strategic reasoning for non-cooperative games, in which decision-makers use certain reasoning algorithms to form probabilistic beliefs about opponents' strategy choices from the commonly known information about opponents' preferences (for an overview, see Perea 2012). For example, the utility proportional beliefs model, developed by Bach and Perea (2014), suggests that real-world decision-makers can be modelled as agents holding utility propor-

tional beliefs – entertaining beliefs on their opponents’ choices proportional to the respective utilities that those choices yield. That is, a player assigns probabilities to opponents’ strategy choices in such a way that, for every opponent, the difference in probability assigned to strategies is proportional to difference in utility that those strategies yield. This assumption is supposed to represent an idea that a real-world decision-maker expects the opponent to choose a strategy yielding a higher personal payoff with a higher probability than a strategy yielding a lower personal payoff, and chooses his or her response accordingly. The players are also assumed to express a common belief in utility-proportional beliefs, which allows them to form beliefs about each other’s beliefs about each other’s actions, thus making the coordination problem resolvable in certain types of games. The theoretical predictions of players’ choices derived from the utility proportional beliefs model seem to better fit with certain experimental findings than the theoretical predictions derived from the standard best-response reasoning model (for details, see Bach and Perea 2014).

An even more sophisticated model of how players form probabilistic beliefs about the opponents’ actions from the available information about opponents’ preferences is Perea’s (2011) proper rationalizability algorithm based on lexicographic belief structures. In Perea’s (2011) model, each player holds a lexicographic belief about his or her opponents’ strategy choices, which is a finite sequence $\lambda_i = (\lambda_i^1, \dots, \lambda_i^K)$ of probability distributions on the set S_{-i} of all the possible combinations of opponents’ strategy choices. A sequence of probability distributions is such that every possible combination $s_{-i} \in S_{-i}$ has a positive probability under some probability distribution λ_i^k in this sequence. For every $k \in \{1, \dots, K\}$, a probability distribution λ_i^k in this sequence is called the level k belief of player $i \in I$. The idea behind the assumption that every possible combination has a positive probability under some probability distribution in the sequence $\lambda_i = (\lambda_i^1, \dots, \lambda_i^K)$ is that each player is ‘cautious’ – does not exclude any combination of opponents’ strategies from consideration, even though s/he believes that some of those combinations are *infinitely more likely* to be played by the opponents than other possible combinations. Player $i \in I$ deems a combination of opponents’ strategies $s_{-i} \in S_{-i}$ infinitely more likely than combination $s'_{-i} \in S_{-i}$ if combination $s_{-i} \in S_{-i}$ has a positive probability under some level k belief λ_i^k , while the combination $s'_{-i} \in S_{-i}$ has zero probability under the first k levels¹¹.

In proper rationalizability model, each player is assumed to respect his or her opponents’ preferences and express a common belief in respect for

¹¹For a technical discussion of the properties of lexicographic probability systems, see Blume et al. 1991a,b.

opponents' preferences – believe that every other player respects his or her opponents' preferences, believe that every other player believes that every other player respects his or her opponents' preferences, and so on *ad infinitum*. A player $i \in I$ is said to express the preferences of player $j_{j \neq i} \in I$ if, for any two j 's strategies $s_j \in S_j$ and $s'_j \in S_j$, player i deems it infinitely more likely that player j will choose s_j than s'_j if player i believes that opponent j prefers strategy s_j over strategy s'_j ¹².

Given the structure of common beliefs, the player performs an iterative addition of restrictions on the set of opponents' combinations of strategies, in light of a belief that every other player performs the same procedure in light of the same belief about every other player, and so on *ad infinitum*. The player starts with a set of all the possible combinations of opponents' strategies S_{-i} and identifies a subset $D_{-i} \subseteq S_{-i}$ of combinations of opponents' strategies, such that each combination of opponents' strategies in the subset D_{-i} is deemed by player i to be infinitely more likely than any combination of opponents' strategies outside the subset D_{-i} . The player continues adding restrictions on the set S_{-i} until no further restrictions can be added. A subset of strategies of player $i \in I$ which can be optimally chosen as responses to combinations of opponents' strategies which survive the iterative addition of restrictions is the set of player i 's properly rationalizable strategies (for an extensive technical discussion and proofs, see Perea 2011).

In principle, lexicographic probability systems are flexible enough to be useful as a reasonable starting point for a formal representation of hypothetical bargainers' beliefs about each other's actions. In Perea's (2011) proper rationalizability model, each player's beliefs about the opponents' strategy choices are assumed to be determined by his or her beliefs about the opponents' preferences over strategies. This assumption, however, can be replaced with a different assumption postulating a different kind of relation between player's beliefs about opponents' choices and the information available to the player about opponents' preferences and/or opponents' beliefs about the structure of the game and its players.

According to Misyak and Chater (2014), hypothetical bargainers choose their actions on the basis of what they would agree to play if they could openly bargain. Hypothetical bargainers identify a combination of strategies which they would agree to play in open bargaining by evaluating the 'goodness' of the feasible agreements and finding an agreement which, given their criterion of 'goodness', they perceive as the 'best' feasible agreement

¹²That is, player $i \in I$ believes that player $j_{j \neq i} \in I$ prefers strategy $s_j \in S_j$ over strategy $s'_j \in S_j$ if there is some level $k \in \{1, \dots, K\}$ such that $u_j(s_j, \lambda_j^k) > u_j(s'_j, \lambda_j^k)$ and $u_j(s_j, \lambda_j^l) = u_j(s'_j, \lambda_j^l)$ for all $l < k$. For details, see Perea 2011.

(Misyak and Chater 2014). In principle, Misyak and Chater’s suggestion can be interpreted as an idea that hypothetical bargainer’s beliefs about the other hypothetical bargainers’ actions are determined by his or her beliefs regarding other hypothetical bargainers’ perceived relative ‘goodness’ of feasible bargains. In other words, a hypothetical bargainer could be assumed to believe that, given the set of available feasible agreements, the opponent, given his or her criterion of ‘goodness’, is *more likely to play a part in the implementation of an agreement which s/he deems better than to play a part in the implementation of an agreement which s/he deems worse*. Assuming that each hypothetical bargainer is using the same criterion of ‘goodness’ and that this fact is common knowledge, it can be shown that hypothetical bargainers could form correlated lexicographic beliefs about each other’s actions.

For an example of how, in principle, hypothetical bargainer’s beliefs could be represented, suppose that two hypothetical bargainers are playing a normal form game $\Gamma = (\{1, 2\}, \{S_i, u_i\}_{i \in \{1, 2\}})$, where S_i is the set of pure strategies and $u_i : \mathbf{S} \rightarrow \mathbb{R}$ is the utility function of player $i \in \{1, 2\}$, where $\mathbf{S} = (S_1 \times S_2)$ is the set of pure strategies of Γ . To simplify this example, I will assume that players only consider pure strategy profiles as possible agreements.

Suppose that every hypothetical bargainer only deems an agreement feasible if it creates no incentive for any of the players to deviate from it. That is, the set of feasible agreements is the set of pure strategy Nash equilibria $\mathbf{S}^{NE} \in \mathcal{P}(\mathbf{S})$. It will be assumed that players are playing a game where $\mathbf{S}^{NE} \neq \emptyset$.

For this particular example, it will be assumed that players evaluate the ‘goodness’ of feasible agreements using a criterion underlying the cardinal BE solution suggested in chapter 3. That is, hypothetical bargainers identify the best feasible agreement by identifying a weakly Pareto optimal Nash equilibrium which minimizes the difference between players’ losses of maximum attainable individual advantage. The choice of the ‘goodness’ criterion in this case is arbitrary, since it will not have any significant effect on the structure of the considered epistemic model.

Let u_i^{max} denote the maximum payoff of player $i \in \{1, 2\}$ associated with some rationalizable outcome of the game, and let u_i^{ref} denote player i ’s reference payoff – the personal payoff, relative to which $i \in \{1, 2\}$ evaluates the individual advantage gains from feasible agreements¹³. For each player $i \in \{1, 2\}$, the level of individual advantage associated with a feasible agree-

¹³For an extensive discussion of how a reference point could be defined, see chapter 3. For the purposes of this discussion, the properties of the reference point will not matter.

ment $\mathbf{s} \in \mathbf{S}^{NE}$ is the following ratio:

$$u_i^t(\mathbf{s}) = \frac{u_i(\mathbf{s}) - u_i^{ref}}{u_i^{max} - u_i^{ref}}. \quad (4.2)$$

It is easy to check that the maximum level of individual advantage attainable to each player relative to his or her reference point is 1. Recall that a benefit-equilibrating solution is a weakly Pareto optimal Nash equilibrium $\mathbf{s} \in \mathbf{S}^{NEwpo}$ which minimizes the difference between players' losses of maximum attainable individual advantage:

$$\phi^c(\mathbf{S}^{NE}) \in \arg \min_{\mathbf{s} \in \mathbf{S}^{NEwpo}} \left\{ \left| (1 - u_i^t(\mathbf{s})) - (1 - u_{j \neq i}^t(\mathbf{s})) \right| \right\}. \quad (4.3)$$

In epistemic terms, hypothetical bargainers can be modelled as players who evaluate the 'goodness' of each weakly Pareto optimal feasible agreement by comparing the difference between players' losses of maximum attainable individual advantage associated with each feasible agreement: An agreement associated with a smaller difference between players' losses of maximum attainable individual advantage is deemed 'better' than the one associated with a larger difference between players' losses of maximum attainable individual advantage. If both players are hypothetical bargainers and this fact is common knowledge, each player can be modelled as holding a belief that, for any two feasible agreements, the opponent is infinitely more likely to play his or her part in realizing an agreement that s/he deems 'better'.

Given the set of agreements $\mathbf{S}^{NEwpo} \subseteq \mathbf{S}^{NE}$, every hypothetical bargainer $i \in \{1, 2\}$ can be modelled as holding a belief that hypothetical bargainer $j_{j \neq i} \in \{1, 2\}$ will play his or her part in realizing some feasible agreement $\mathbf{s} \in \mathbf{S}^{NEwpo}$ by playing his or her *component* in the strategy profile $\mathbf{s} = (s_1, s_2)$. Player j 's component in strategy profile $\mathbf{s} \in \mathbf{S}^{NEwpo}$ will be denoted as $c_j(\mathbf{s})$. Given the set of agreements \mathbf{S}^{NEwpo} , player $i \in \{1, 2\}$ can identify the set of components $C_j \subseteq S_j$ that player j would choose in order to play a part in the implementation of some agreement from the set \mathbf{S}^{NEwpo} .

An epistemic model of a game is a tuple $\mathcal{M} = (T_i, \lambda_i)_{i \in \{1, 2\}}$, where T_i is the set of possible types of player $i \in \{1, 2\}$ and λ_i is a function which assigns a lexicographic probability system $\lambda_i(t_i)$ to every type $t_i \in T_i$. A lexicographic probability system $\lambda_i(t_i)$ represents the beliefs that type $t_i \in T_i$ holds about the *component* of the strategy profile that his or her opponent $j_{j \neq i} \in \{1, 2\}$ is going to play in order to implement some agreement in \mathbf{S}^{NEwpo} , as well as beliefs about j 's beliefs (i.e. j 's type). A lexicographic belief $\lambda_i(t_i) = (\lambda_i^1, \dots, \lambda_i^K)$ of player $i \in \{1, 2\}$ is a lexicographic probability system on the set of possible states $C_{j \neq i} \times T_j$, which is the set of all the possible component-type combinations of player $j \in \{1, 2\}$. Every component-type

combination $(c_j(\mathbf{s}), t_j) \in C_j \times T_j$ can be interpreted as i 's belief that player $j \in \{1, 2\}$ plays his or her component of the strategy profile $\mathbf{s} \in \mathbf{S}^{NE_{wpo}}$ and holds a lexicographic belief $\lambda_j(t_j)$. For every component–type combination $(c_j(\mathbf{s}), t_j) \in C_j \times T_j$, we can define its rank $\wp((c_j(\mathbf{s}), t_j), \lambda_i(t_i))$ within lexicographic probability system $\lambda_i(t_i)$ of type $t_i \in T_i$ as the *lowest* level k such that $\lambda_i^k(c_j(\mathbf{s}), t_j) > 0$. Type $t_i \in T_i$ is said to deem component–type combination $(c_j(\mathbf{s}), t_j)$ infinitely more likely than component–type combination $(c_j(\mathbf{s}'), t'_j)$ if and only if

$$\wp((c_j(\mathbf{s}), t_j), \lambda_i(t_i)) < \wp((c_j(\mathbf{s}'), t'_j), \lambda_i(t_i)). \quad (4.4)$$

If the fact that players are hypothetical bargainers is common knowledge among the interacting players, then every type $t_i \in T_i$ knows that every type $t_j \in T_j$ is a hypothetical bargainer, and so deems the agreement $\mathbf{s} \in \mathbf{S}^{NE_{wpo}}$ ‘better’ than the agreement $\mathbf{s}' \in \mathbf{S}^{NE_{wpo}}$ only if

$$\left| (1 - u_i^t(\mathbf{s})) - (1 - u_{j \neq i}^t(\mathbf{s})) \right| < \left| (1 - u_i^t(\mathbf{s}')) - (1 - u_{j \neq i}^t(\mathbf{s}')) \right|. \quad (4.5)$$

Given the assumptions about the relationship between i 's beliefs about the relationship between j 's actions and j 's evaluations of outcomes, as well as i 's beliefs about j 's evaluations of feasible agreements, the following relationship will hold for every type $t_i \in T_i$:

$$\wp((c_j(\mathbf{s}), t_j), \lambda_i(t_i)) < \wp((c_j(\mathbf{s}'), t_j), \lambda_i(t_i)) \quad (4.6)$$

only if the following condition is satisfied:

$$\left| (1 - u_i^t(\mathbf{s})) - (1 - u_{j \neq i}^t(\mathbf{s})) \right| < \left| (1 - u_i^t(\mathbf{s}')) - (1 - u_{j \neq i}^t(\mathbf{s}')) \right|. \quad (4.7)$$

Given the aforementioned assumptions, each player will have a lexicographic ranking of opponent's actions, in which opponent's choice of a component of an agreement with a smaller difference between players' losses of maximum attainable individual advantage will have a lower rank than the choice of a component of an agreement with a larger difference between players' losses of maximum attainable individual advantage. In other words, each player will hold a belief that his or her opponent is infinitely more likely to play his or her component of an agreement with a smaller difference between players' losses of maximum attainable individual advantage than his or her component of an agreement with a larger difference between players' losses of maximum attainable individual advantage. The feasible agreement with the lowest rank will be the one which, given the set of agreements,

minimizes the difference between players' losses of maximum attainable individual advantage. If the fact that both players are hypothetical bargainers is common knowledge, each player knows that the other player holds the same lexicographic belief, knows that the other player knows this, and so on *ad infinitum*.

Suppose that type $t_i \in T_i$ holds a lexicographic belief $\lambda_i(t_i) = (\lambda_i^1, \dots, \lambda_i^K)$ on the set $C_{j \neq i} \times T_j$. For every level $k \in \{1, \dots, K\}$, the expected utility from playing a component $c_i(\mathbf{s})$ of agreement $\mathbf{s} \in \mathbf{S}^{NEwpo}$ can be defined as follows:

$$u_i(c_i(\mathbf{s}), \lambda_i^k) = \sum_{(c_j(\mathbf{s}), t_j) \in C_j \times T_j} \lambda_i^k(c_j(\mathbf{s}), t_j) u_i(c_i(\mathbf{s}), c_j(\mathbf{s})). \quad (4.8)$$

Recall that a component–type combination with a lower rank in lexicographic probability system is deemed infinitely more likely than any component–type combination with a higher rank. This means that, for every type $t_i \in T_i$, it is always the case that, for any two agreements $\mathbf{s} \in \mathbf{S}^{NEwpo}$ and $\mathbf{s}' \in \mathbf{S}^{NEwpo}$,

$$u_i(c_i(\mathbf{s}), \lambda_i(t_i)) > u_i(c_i(\mathbf{s}'), \lambda_i(t_i)) \quad (4.9)$$

if it is the case that

$$\wp((c_j(\mathbf{s}), t_j), \lambda_i(t_i)) < \wp((c_j(\mathbf{s}'), t_j), \lambda_i(t_i)). \quad (4.10)$$

Since the agreement with the lowest rank is the one which minimizes the difference between players' losses of maximum attainable individual advantage, a hypothetical bargainer will be motivated to play his or her component of that agreement. If it is common knowledge that each player is a hypothetical bargainer, each player knows that the other player holds the same lexicographic belief and is motivated to play his or her component of the same agreement, knows that the other player knows this, and so on *ad infinitum*. Thus, given the aforementioned assumptions about the relation between player's beliefs about opponent's actions and opponent's evaluations of outcomes, hypothetical bargainers will always be motivated to play their part in the attainment of an outcome that they identify as the BE solution of Γ .

It is important to stress that the considered model of lexicographic beliefs should not be viewed as a model offering a compelling answer to the question of why hypothetical bargainers are motivated to play their part in the implementation of the bargaining solution, but merely as a theoretical exploration of how, *in principle*, hypothetical bargaining could be modelled as a belief–formation process. The considered model has considerable empirical and conceptual limitations. Although lexicographic beliefs become increasingly popular in epistemic game theory, there is no evidence that real-world

decision-makers can form beliefs, the structure of which even remotely resembles the structure of lexicographic beliefs. Another major conceptual problem is that lexicographic probability systems can be used to rationalize virtually any strategy choice, provided that the strategy is rationalizable in the standard game theoretic sense and that no restrictions are imposed on how players use the information about the game and its players in forming beliefs about the opponents' actions.

The assumed relationship between player's beliefs about the opponents' actions and player's beliefs about the opponents' evaluations of outcomes may also be viewed as conceptually problematic: Although each player's probabilistic beliefs about the possible states of the world (i.e. combinations of opponents' strategies) are independent from player's own evaluations of outcomes, the player is assumed to believe that the probability of each possible state is not independent from opponents' evaluations of outcomes. That is, each player believes that the opponent is more likely to play a part in the attainment of an outcome that s/he personally deems better than to play a part in the attainment of an outcome which s/he personally deems worse. This assumption can be criticized as being a departure from the standard rational choice theory, since it implies that each player believes that his or her opponents violate the so-called *Aesop's principle*, which requires rational decision-maker's preferences and beliefs to be independent of each other (for extensive discussion, see, for example, Binmore 2009a and Perea 2012).

Nevertheless, an interpretation of hypothetical bargaining which treats it as a belief-formation algorithm seems to involve a less radical departure from the rational choice framework than any interpretation which treats it as a choice algorithm. A mechanistic interpretation of hypothetical bargaining must be based on assumption that hypothetical bargainers always play their part in the attainment of an outcome which they recognize as a bargaining solution, irrespective of what they believe about the other players' actions. This interpretation implies that a hypothetical bargainer will try to achieve the identified goal – implement the hypothetical bargaining agreement – in situations where s/he has no reason to believe that the other players will do their part in the attainment of that goal.

If, on the other hand, hypothetical bargaining were to be interpreted as a belief-formation algorithm, then each hypothetical bargainer's choice could at least be shown to be consistent with his or her beliefs about the choices of other hypothetical bargainers. As has been pointed out by Ulcina and Orbano (1994), the standard best-response reasoning model is not a belief formation model but rather a choice consistency model: A rational player is required to choose a strategy which is optimal in light of his or her beliefs

about the opponents' actions. The theory, however, does not provide an answer to the question of how rational players should form beliefs about opponents' rationalizable strategy choices (Olcina and Urbano 1994).

The proponents of hypothetical bargaining suggest that this theory should be viewed as a model of rational *choice*. I have argued that their suggested interpretation of the hypothetical bargaining theory is not tenable without an adequate theory of why and under what conditions a hypothetical bargainer's decision to play a part in the attainment of an outcome identified as the bargaining solution of a game is a rational choice. Until an adequate explanation is provided, the model remains incomplete. If hypothetical bargaining were to be interpreted as a belief-formation model, it could at least be shown to be consistent with the most basic choice consistency principle of the standard game theory: Hypothetical bargainers could be modelled as players who choose a best-response to their beliefs about opponents' strategy choices. However, it seems that such a belief-formation model may not be possible without a non-trivial departure from the epistemic rationality principles endorsed by the orthodox game theory, which means that hypothetical bargaining, *if interpreted as a model of rational strategy choice rather than merely as a model of how players identify mutually beneficial solutions of games*, may not be compatible with the epistemic principles of orthodox game theory.

4.3 Hypothetical Bargaining and Social Coordination

4.3.1 Hypothetical Bargaining and the Problem of Common Beliefs

Besides the aforementioned failure of the theory to explain the reason of why hypothetical bargainers would have a rational incentive to play their part in the attainment of outcomes identified as bargaining solutions, a rational choice interpretation of the theory also makes it vulnerable to the problem of common beliefs. Even if hypothetical bargainers were always motivated to play their part in the attainment of the identified bargaining solution under common knowledge assumptions, the descriptive relevance of the theory can still be questioned on empirical grounds.

In real world social interactions, social agents face uncertainty about each other's type. A social agent who reasons as a hypothetical bargainer may thus be uncertain as to whether the other player is a hypothetical bargainer

or not. In that case, hypothetical bargainer's motivation to play a part in the attainment of an outcome that s/he identifies as the bargaining solution of a particular interdependent decision problem will depend on his or her beliefs about the other decision-maker's type.

For example, consider the three strategy coordination game depicted in Figure 4.2. This game has three pure strategy Nash equilibria: $(u, l), (m, c)$

	l	c	r
u	7, 3	0, 0	0, 2
m	0, 0	5, 5	0, 2
d	2, 0	2, 0	3, 7

Figure 4.2: Coordination game

and (d, r) . There are also two Nash equilibria in mixed strategies¹⁴. Suppose that the row player is a hypothetical bargainer who identifies the Nash equilibrium (m, c) as the unique hypothetical bargaining solution of this game. The row player believes that if the column player is a hypothetical bargainer, s/he will play c with probability 1. However, the row player is uncertain whether the column player is a hypothetical bargainer, and so assigns probability $p > 0$ to the event of the column player being of another type than s/he is.

Suppose that the row player believes that there is some probability $p > 0$ that the column player is a type of player who plays each pure strategy at random (that is, plays each strategy with probability $1/3$). The row player thus believes that the probability of the column player being a hypothetical bargainer is $(1 - p)$. If the opponent were a hypothetical bargainer, the row player's best response would be strategy m . If, on the other hand, the opponent were to choose each of the pure strategies at random, the optimal response for the row player would be to play either strategy u or strategy d (notice that both strategies would yield the row player the same expected payoff of $2\frac{1}{3}$). It is easy to check that the row player is indifferent between playing strategies u, d and m when $p = 15/17$. If the row player were to believe that the probability of the column player being a hypothetical bargainer were higher than $2/17$, s/he would be motivated to play a part in the attainment of the outcome (m, c) .

¹⁴The two mixed strategy Nash equilibria of this game are: $(\frac{7}{8}u, \frac{1}{8}d; \frac{3}{8}l, \frac{5}{8}r)$, $(\frac{7}{10}m, \frac{3}{10}d; \frac{1}{2}c, \frac{1}{2}r)$.

The result would be different if the row player were to hold a different belief regarding the actions of non-hypothetical bargainer. Suppose, for example, that the row player believes that the column player may be a type of decision-maker who always chooses his or her *maximin* strategy. The row player thus believes that the column player may choose his or her *maximin* strategy r . If the column player chooses strategy r , the row player's best-response is to play d . It is easy to check that the row player is indifferent between playing strategies d and m when $p = 5/8$. It means that the row player would be motivated to play his or her part in the attainment of outcome (m, c) if the probability of the column player being a hypothetical bargainer were higher than $3/8$.

This example shows that for a hypothetical bargainer to be motivated to play a part in the attainment of an outcome which s/he identifies as the hypothetical bargaining solution, s/he has to believe that the probability of the opponent being a hypothetical bargainer is sufficiently high. An answer to the question of whether players' beliefs must be common knowledge in order to sustain their motivation to act as hypothetical bargainers depends on how the model of hypothetical bargaining is interpreted. If a hypothetical bargainer is assumed to be a decision-maker who always believes that every other hypothetical bargainer always plays his or her part in the attainment of the identified bargaining solution, then hypothetical bargainers' beliefs about each other's beliefs do not need to be common knowledge. If, on the other hand, a hypothetical bargainer is assumed to be a decision-maker who believes that other hypothetical bargainers' motivation to play the appropriate strategies is *conditional* on their belief that the probability of their opponents being hypothetical bargainers is sufficiently high, then a hypothetical bargainer must not only believe that the probability of the opponents being hypothetical bargainers is sufficiently high, but also believe that each hypothetical bargainer believes this, believe that each hypothetical bargainer believes that each hypothetical bargainer believes this, and so on *ad infinitum*¹⁵.

More specifically, hypothetical bargaining can operate in a population P if there is a subset $P^{hb} \subseteq P$ of hypothetical bargainers, such that each

¹⁵According to Gintis 2008, the game theoretic models which purport to explain individual's decision to comply with a social norm as a rational choice face a somewhat similar problem: For a rational individual to be motivated to comply with a norm in a situation of type C , s/he has to believe that the other rational individuals are aware that the norm applies to situations of type C and will actually follow it. Each rational individual's compliance with the norm is thus conditional on expectation that other rational individuals will comply with the norm as well. Therefore, each player must believe that every other player will comply with the norm, believe that every other player believes this, and so on.

individual $i \in P^{hb}$ believes that:

1. There exists a subset P^{hb} of hypothetical bargainers in the population P ,
2. The subset $P^{hb} \subseteq P$ is sufficiently large,
3. Each individual $j_{j \neq i} \in P^{hb}$ believes that 1 and 2 is the case,
4. Each individual $j_{j \neq i} \in P^{hb}$ believes that each individual $k_{k \neq j} \in P^{hb}$ believes that 1 and 2 is the case,
5. Each individual $j_{j \neq i} \in P^{hb}$ believes that each individual $k_{k \neq j} \in P^{hb}$ believes that each individual $l_{l \neq k} \in P^{hb}$ believes that 1, 2 is the case,

and so on¹⁶.

In other words, hypothetical bargaining can operate as a coordination device in situations where hypothetical bargainers express a common belief that the proportion of hypothetical bargainers in the population is sufficiently large. Given this necessary (due to rationalization problem, possibly not sufficient) epistemic requirement, the question concerning the conditions under which hypothetical bargainers could be expected to express such a common belief becomes particularly important.

Misyak et al. (2014) argue that although hypothetical bargaining relies on common knowledge assumptions, this epistemic requirement is not particularly problematic, since ‘a history of social interactions, communication and common culture can foster common knowledge among individuals, thereby facilitating virtual bargaining’, while ‘previous virtual bargains (and, more broadly, real bargains and past outcomes) provide precedents for current and future virtual bargains’ (Misyak et al. 2014: 516). This explanation seems to be in line with Gintis’s (2008) suggestion that common beliefs emerge

¹⁶Note that the structure of common beliefs is relatively simple due to assumption that each hypothetical bargainer believes that every hypothetical bargainer always plays his or her part in the attainment of an outcome that s/he identifies as the bargaining solution of a game. In other words, it is assumed that the rationalization problem discussed in section 4.2 does not arise in the first place. Otherwise the structure of beliefs would be much more complicated, since each hypothetical bargainer’s motivation to play a part in the implementation of the hypothetical bargaining solution would be conditional on a belief that other hypothetical bargainers will play their part in the attainment of that outcome. This would obviously make hypothetical bargaining an even more problematic social coordination theory.

as a result of common experiences of social agents. The idea is that individuals who live in the same population form similar beliefs by repeatedly interacting with other individuals from the same population. That is, by repeatedly interacting with other individuals from the same population, an individual begins to recognize the differences in other individuals' behaviour and, provided that communication is possible, even differences in how other individuals conceptualize the decision problems. At least some individuals begin to classify other decision-makers into types. An individual who develops the ability to classify decision-makers into types can also learn to distinguish the types of decision-makers that s/he interacts with more frequently from those types which are less prevalent in his or her environment. Such an individual can thus form a belief about the relative frequency of each type of decision-maker in his or her environment. Since individual knows that his or her beliefs are based on the experience gained through repeated interactions with other individuals, s/he expects the other individuals from the same population to have gained similar experience and, consequently, to have similar beliefs about the relative frequencies of decision-makers' types as s/he does. As is explained by Gintis, 'the members of our species, *H. Sapiens*, have the capacity to conceive that other members have minds and respond to experience in a manner parallel to themselves' (Gintis 2008: 140). Therefore, 'if agent i believes something, and if i knows that he shares certain environmental experiences with agent j , then i knows that j probably believes this thing as well' (Gintis 2008: 140).

Even if we were to accept Gintis's story as a plausible explanation of how common beliefs emerge (for a critical discussion, see Binmore 2008), it relies on assumption that hypothetical bargainers interact in a population where the proportion of hypothetical bargainers is already sufficiently large. In other words, for a hypothetical bargainer to form a belief that the proportion of hypothetical bargainers in the population is sufficiently large, s/he must be interacting with other hypothetical bargainers sufficiently frequently. This would be very unlikely to happen in a population with very few hypothetical bargainers and a large number of individuals using other decision-making approaches. The theory seems to fall into a vicious cycle: Hypothetical bargainers can only form a belief which motivates them to act by interacting in a population where a sufficiently large proportion of the population are acting as hypothetical bargainers, yet such a population state is only present when a sufficient number of players are already acting as hypothetical bargainers. In other words, the theory cannot explain how a population can reach a state where the number of hypothetical bargainers is large enough to sustain hypothetical bargainers' motivation to act as

hypothetical bargainers.

There seems to be a plausible response to this problem which, however, involves a departure from the rational choice interpretation of hypothetical bargaining. Hypothetical bargaining could be viewed as a reasoning algorithm which have evolved out of a relatively primitive choice *heuristic* – a simple decision rule which required little to no strategic deliberation¹⁷ and which, due to its fitness-enhancing properties, was able to spread in the population either via genetic inheritance or via cultural imitation.

For this explanation to be credible, a primitive version of hypothetical bargaining must be shown to be an evolutionary advantageous decision rule. This possibility will be considered in the next section.

4.4 Hypothetical Bargaining and Evolution

4.4.1 Evolutionary Game Theory and Population Dynamics

Historical research does not offer much in terms of showing that certain decision rules were more advantageous than others. In order to fill this explanatory gap, game theorists and other social scientists have turned to evolutionary game theoretic models of social interactions. Relatively simple evolutionary game theoretic models can be used to explore the possibility that individuals who were making choices *as if* they were using a certain decision rule would, given certain initial population conditions, be evolutionary successful in certain types of social interactions, since their responses would have given them a fitness advantage over individuals who were making choices *as if* they were using other decision rules.

In simple evolutionary models, such as in those based on replicator dynamics (Taylor and Jonker 1978), each player is assumed to be playing one of the pure strategies of the game. It is also assumed that there is a certain small probability $\epsilon > 0$ that each player will ‘mutate’ – spontaneously change his or her strategy. In each ‘round’ of an evolutionary game, an individual plays against an opponent from the same population (in simplest models, the players are assumed to be drawn from the population at random). The population consists of a large but finite number of individuals, each of which is programmed to play some pure strategy s^i of the game, where $i \in \{1, \dots, K\}$ is the set of pure strategies of an evolutionary game. Letting $p_i(t) \geq 0$ denote

¹⁷For extensive discussion of decision-making heuristics, see, for example, Pearl 1983, Gigerenzer et al. 1999, Gigerenzer and Selten 2001 and Alexander 2007.

the number of players using pure strategy s^i at time t and $p(t) = \sum_{i \in K} p_i(t)$ denote the total population at time t , we can define the *population share* playing strategy s^i at time t as follows:

$$x_i(t) = \frac{p_i(t)}{p(t)}. \quad (4.11)$$

At every point in time t , a population can be characterized by its state – the population shares playing each of the pure strategies of the game. That is, a *population state* is a vector $x(t) = (x_1(t), \dots, x_k(t))$, where each component $x_i(t)$ is the population share using strategy s^i at time t . When population shares playing each of the pure strategies of the game change, a population state is said to change. Notice that each population state is formally identical to a mixed strategy¹⁸.

Assuming that players are matched at random, the expected payoff associated with a pure strategy s^i given the population state $x(t)$ is $u(s^i, x)$. The *population average payoff* – the expected payoff of an individual drawn from the population at random – can be defined as follows:

$$u(x, x) = \sum_{i=1}^k x_i u(s^i, x). \quad (4.12)$$

In evolutionary models based on replicator dynamics, individual's expected payoff is interpreted as *fitness* – the expected number of offspring or imitators¹⁹ who will be using individual's strategy. Those individuals whose strategies earn them a higher expected payoff than the population average

¹⁸That is, a vector representing a population state can be interpreted as a mixed population strategy, where each population share is the probability weight assigned to a pure strategy of the game. Therefore, each player can be viewed as using his or her pure strategy against the mixed strategy of the population. For details, see Thomas 1984.

¹⁹Replicator dynamics, introduced by Taylor and Jonker (1978) is a simple and arguably the most widely-used evolutionary dynamics model. It is an abstract mathematical model which can be interpreted as representing either biological or cultural evolution. If the model is interpreted as a model of biological evolution, individual's payoff is interpreted as the expected number of offspring that will inherit a gene or genes responsible for producing player's phenotypic trait, such as a certain behavioural response to a specific environmental stimulus. If the model is treated as a model of cultural evolution, individual's payoff is interpreted as the expected number of players who will imitate player's behaviour. Evolutionary learning models, such as those based on best response dynamics, are specifically tailored to represent the dynamics of a population of decision-makers who can consciously change their strategy, and so are more suitable for representing cultural rather than biological evolution. For extensive technical discussion of the different evolutionary models and their assumptions, see Weibull 1995, Fudenberg and Levine 1998, Alexander 2007 and Sandholm 2010a.

payoff tend to have more offspring or imitators than the population average. Therefore, the proportion of players who use such evolutionary successful strategies tends to increase over time.

Assuming that the reproduction (imitation) is continuous over time, and that individuals' background fitness $\beta \geq 0$ is independent from the outcomes of the evolutionary game in question, the birth/imitation rate of an individual using strategy s^i can be defined as follows:

$$\{\beta + u[s^i, x(t)]\}. \quad (4.13)$$

The following discussion will rely on a standard assumption that the death rate $\delta \geq 0$ is the same for all individuals. With dots for time derivatives and suppressing the time arguments, the replicator dynamics can be defined as follows:

$$\dot{p}_i = [\beta + u(s^i, x) - \delta] p_i. \quad (4.14)$$

The corresponding replicator dynamics for each population share x_i can be defined as follows:

$$\dot{x}_i = [u(s^i, x) - u(x, x)] x_i. \quad (4.15)$$

A population is said to be in a *stationary, or rest, state* when no pure strategy earns its user a higher expected payoff than the population average payoff. In some of the stationary states, virtually everyone in the population uses the same pure strategy, while in other stationary states two or more pure strategies may be used by substantial proportions of the population²⁰.

In evolutionary models based on replicator dynamics, players are modelled as individuals programmed to play a specific pure strategy of the game. That is, individuals do not update their strategies, even if their strategies yield poor payoffs. Evolutionary models based on other types of dynamics can, however, be used to model the learning process of individuals capable of consciously adopting other individuals' strategies in response to information that other individuals' strategies yield a better expected payoff than their own strategies. For example, the *best response dynamics* (Gilboa and Matsui 1991, Matsui 1992) can be used to model the learning of myopic best responders. A myopic best responder is a decision-maker who adopts the strategy that will confer the highest expected payoff in the next generation under assumption that other individuals in the population will not change their strategies. If there is more than one such strategy, a myopic best responder chooses one at random.

²⁰For an in-depth technical discussion of replicator dynamics, see Taylor and Jonker 1987, Weibull 1995, Hofbauer and Sigmund 1998 and Sandholm 2010a.

A myopic best responder is a decision-maker similar to the best-response reasoner insofar as s/he performs expected payoff calculations and chooses a strategy yielding the highest expected payoff. However, a best responder is not a Bayesian-rational decision-maker²¹, since s/he does not engage in complex deliberations to figure out an optimal response to the expected *future* population state – a population state that a Bayesian-rational decision-maker expects to encounter after other individuals will update their strategies²². Compared to a perfectly rational Bayesian deliberator, a myopic best-response reasoner is a boundedly rational agent (for extensive discussion, see Alexander 2007).

The best response dynamics seems to offer a more realistic representation of how real-world decision-makers update their strategies than either the replicator dynamics or the Bayesian models. People are capable of learning from their own mistakes, as well as from the mistakes and successes of others, yet at the same time they are not Bayesian-rational decision-makers with unlimited computational capabilities. In evolutionary models based on best response dynamics, individuals update their strategies, yet they do that purely on the basis of the observed results of other individuals' actions, not on the basis of complex calculations of what the other individuals can be expected to do in the future. This model of learning captures the idea that social agents improve their own actions by observing what the other social agents do in certain situations, as well as by comparing the success of their own actions to the success of the actions of other social agents (for extensive discussion, see Bicchieri 2006). In addition, the relatively problematic common knowledge assumptions underlying the Bayesian learning models – common knowledge of rationality and (often) common knowledge of priors – do not play any role in evolutionary learning models based on best response dynamics (for extensive discussion, see Nisan et al. 2011).

Formally, best response dynamics is based on assumption that every player in the population is a myopic best-response reasoner who, when given a chance to update his or her strategy, does it by following a best response

²¹A Bayesian rational player is usually assumed to be a decision-maker who always uses all the available information to form consistent beliefs about the opponents' choices and has no computational limitations. Such a player takes all the information about the payoffs received by other players, and, assuming that s/he knows how they update their strategies in response to received payoffs, calculates the changes in the proportions of players using the pure strategies available in the game. A Bayesian rational player then chooses a strategy which is optimal in light of his or her beliefs about the future population state.

²²In other words, a myopic best-responder is assumed to be incapable to predict the future mixed strategy of the population, and so s/he cannot choose a best-response to it.

rule²³.

Suppose that $\mathcal{BR}(x)$ is the set of all (mixed) best responses to a population state (i.e. mixed population strategy) x ²⁴. With dots for time derivatives and suppressing time arguments, the best response dynamics can be representing as follows:

$$\dot{x} = \mathcal{BR}(x) - x. \quad (4.16)$$

It is easy to see that best response dynamics is the classical Cournot best response process, which in the present context would read

$$x(t+1) = \mathcal{BR}[x(t)]^{25}. \quad (4.17)$$

When applied in the analysis of decision rules, evolutionary game theoretic approach can reveal two things. First, it can show that a decision rule has ‘resistance’ against other decision rules. That is, it can show that a stationary population state where everyone is using the same decision rule to choose a pure strategy of the game is an *evolutionarily stable state* – a population state, such that no ‘mutant’ using a different decision rule leading to a choice of a different pure strategy or strategies can, in evolutionary time, invade and take over the population, provided that the share of mutants emerging in the population at the same time is not too large²⁶. If the population is in an evolutionarily stable state where everyone is using

²³The best response dynamics is similar to the deliberational dynamics of rational myopic players suggested by Skyrms (1990). The major difference is that best response dynamics is a population model, and so the state space represents all the possible population states, while in Skyrms’s model the state space represents all the possible combinations of players’ subjective probability distributions over pure strategies. For details, see Skyrms 1990.

²⁴In some games, the set $BR(x)$ is not a singleton, meaning that multiple trajectories can sprout from a single initial state, and solution trajectories can cycle in and out of the Nash equilibria. In such cases, the dynamics can be interpreted as a differential inclusion: $\dot{x} \in BR(x) - x$ (Hofbauer 1995b). This problem can be avoided with the introduction of random perturbations of player’s payoffs, meaning that players called to adjust their strategy would not always choose a best response (Hopkins 1999). In some cases, however, the smooth best response dynamics can only be achieved with the introduction of large perturbations of players’ payoffs, and so it is questionable whether such models can be interpreted as representing the behaviour of best responders. For a detailed discussion of this problem and possible solutions, see Hofbauer 1995b (unpublished manuscript), Hopkins 1999, and Sandholm 2010a.

²⁵For a detailed technical discussion, see Gilboa and Matsui 1991, Matsui 1992, Hofbauer 1995a,b, Hopkins 1999 and Sandholm 2010a.

²⁶Each evolutionarily stable state has an *invasion barrier*, which is the maximally large proportion of mutants entering the population at the same time that an evolutionarily stable state can resist. For a technical discussion, see Hofbauer and Sigmund 1998, Weibull 1995.

the same decision rule that leads to the choice of the same strategy, then the strategy prescribed by decision rule is an *evolutionarily stable strategy* – a strategy which, if used by every individual in the population, makes the population uninvadable by mutants using other strategies, provided that the initial proportion of invading mutants is not too large (see, for example, Maynard Smith and Price 1973)²⁷.

Suppose that a population is in a state where every individual is using an *incumbent* decision rule which always leads him or her to playing strategy x , and so each individual's payoff from any interaction with another player can be defined as $u(x, x)$. Suppose that ε is the share of mutants using a decision rule which leads them to playing strategy y ²⁸, where $\varepsilon \in (0, 1)$. Assuming that player's interactions are random, the probability of each individual playing against the mutant is ε , while the probability of playing against the individual using the incumbent decision rule is $(1 - \varepsilon)$. Individual's expected payoff from using an incumbent decision rule in such a mixed population state can be defined as follows:

$$u(x, \varepsilon y + (1 - \varepsilon)x). \quad (4.18)$$

Strategy x is said to be evolutionarily stable if and only if

$$u(x, \varepsilon y + (1 - \varepsilon)x) > u(y, \varepsilon y + (1 - \varepsilon)x) \forall y \neq x. \quad (4.19)$$

An equivalent way of stating that strategy x is an evolutionarily stable strategy is to say that it satisfies the first-order (equation 4.20) and the second-order (equation 4.21) best response conditions²⁹:

²⁷Technically, a population is in an evolutionarily stable state if, after a disturbance, it returns to playing a certain strategy or mix of strategies, provided that the disturbance is not too large. Under replicator dynamics, a population state which is asymptotically stable is an evolutionarily stable state (Taylor and Jonker 1978, Hines 1980). It has been shown that every regular evolutionarily stable strategy is asymptotically stable under the general imitative dynamics (Cressman 1997), any impartial pairwise comparison dynamics, such as the Smith dynamic, any separable excess payoff dynamics, such as the Brown-von Neumann-Nash dynamic, and under the best response dynamics (Sandholm 2010b).

²⁸In evolutionary game theory models, the set of possible mutant strategies is often assumed to be the set of pure strategies available in the game. In other words, it is assumed that no mutant playing exogenous strategy (i.e. strategy which is not available in the original game) is possible. This assumption simplifies the tractability and analysis of the model, but may be seen as failing to account for the exogenous shocks that may radically change the structure of the population game, and, consequently, the evolutionary dynamics.

²⁹For a technical discussion of why this is the case, see Weibull 1995.

$$u(y, x) \leq u(x, x) \forall y, \quad (4.20)$$

$$u(y, x) = u(x, x) \Rightarrow u(y, y) < u(x, y) \forall y \neq x. \quad (4.21)$$

It is important to note that an evolutionarily stable state need not be a state where everyone plays the same evolutionarily stable strategy. In *polymorphic* evolutionarily stable states, several pure strategies are played by substantial population shares (for extensive discussion, see Weibull 1995 and Skyrms 1996).

Second, evolutionary game theoretic analysis can reveal the conditions under which a certain decision rule could spread in the population. Evolutionary dynamics is always modelled as a process starting from some initial population state, which is some vector of population shares playing different strategies³⁰. A decision rule which picks a strategy that, starting with the initial population state, is evolutionary successful may spread in the population up to a point where an absolute majority of individuals will be using that strategy. If a population state where everyone is using that strategy is an evolutionarily stable state, the evolutionary model can be viewed as offering a plausible story of how a decision rule could have ‘taken over’ the population, provided, of course, that the initial population state from which such a takeover can be shown to be theoretically possible is a *plausible representation of the real-world population conditions*.

The question of which initial population states represent the plausible real-world population conditions does not have a clear answer. Evolutionary models are highly idealized, and the precise criteria of how the realism of the parameter values of such models ought to be evaluated has no satisfactory answer. There seems to be a general rule that a scenario in which a strategy can be shown to spread successfully when the population share using it is small in the initial population state is a more plausible evolutionary story than the one where a strategy can only be shown to spread successfully if the population share using it in the initial population state is large³¹. This

³⁰In some initial population states, one or more pure strategies may be completely absent, which means that the proportion of players using the absent strategies is 0. This, however, does not mean that such strategies will be absent for the whole duration of the evolutionary process: Assuming that spontaneous mutations are possible, a mutant playing a strategy which was absent in the initial population state can appear in later stages of the process.

³¹For an extensive discussion of the explanatory significance of the initial population conditions in evolutionary game theoretic models, see, for example, Skyrms 1996 and Alexander 2007.

criterion is mostly based on intuition that a formation of a population state in which a population share using a particular decision rule is small seems to be a more likely event than the formation of a population state with a large population share using that rule. This intuition is, to some extent, supported by the theoretical assumptions about the spontaneous mutations: The probability of a small number of individuals *simultaneously* switching to using the same decision rule due to independent random mutations is higher than the probability of a large number of individuals mutating in the same way simultaneously³².

4.4.2 Hypothetical Bargaining as an Evolutionary Adaptation

Assuming that hypothetical bargainers always play a (weakly) Pareto optimal Nash equilibrium, a population where everyone is using such a decision rule will be in an evolutionarily stable state if the bargaining solution of a game is a strict Nash equilibrium: A population state in which everyone is playing a strict Nash equilibrium of the game is always evolutionarily stable (for a technical discussion and proofs, see van Damme 1987 and Weibull 1995). This result, however, does not give much evolutionary justification for hypothetical bargaining. All it suggests is that a population where everyone is a hypothetical bargainer and hypothetical bargaining solution of a game is a strict Nash equilibrium cannot be invaded by mutants using decision rule which prescribes any other strategy of the game, provided that the share of invading mutants is not too large. However, any mutant decision rule which prescribes its users the same strict Nash equilibrium as the hypothetical bargaining can invade the population of hypothetical bargainers and survive in it. In addition, in games where hypothetical bargaining solution is not a strict Nash equilibrium, a population of hypothetical bargainers may not

³²Suppose a finite population where every player is playing strategy s^i . There are two possible mutant strategies that could invade the population – strategies s^j and s^k . Suppose that each player's mutation rate (i.e. mutation probability) is $\epsilon \in (0, 1)$, and that mutation events are independent. Suppose that the probability of a mutating player switching from strategy s^i to either strategy s^j or strategy s^k is $\frac{1}{2}$. Thus, the probability of each mutating player ending up using strategy s^j or strategy s^k is the same – $\frac{\epsilon}{2}$. Suppose that, for the strategy s^j to exceed the invasion threshold (that is, to spread in the population), 7 individuals must *simultaneously* switch their strategies from s^i to s^j , while 9 individuals have to simultaneously switch from strategy s^i to strategy s^k in order for strategy s^k to spread. The probability of 7 players simultaneously switching from strategy s^i to strategy s^j is $\left(\frac{\epsilon}{2}\right)^7$, while the probability of 9 players simultaneously switching from strategy s^i to strategy s^k is $\left(\frac{\epsilon}{2}\right)^9$. The probability of the former event is higher than that of the latter.

resist the invasion of mutants, since it may not be in an evolutionarily stable state.

This means that hypothetical bargaining, if interpreted as a decision rule leading players to playing a Nash equilibrium of a game, will be resistant to invasions of mutants using decision rules which lead them to playing other strategies than strategies used by hypothetical bargainers, but this will be true only in games where hypothetical bargaining solution is a strict Nash equilibrium. This result does not rule out the possibility that hypothetical bargaining could have survived in the evolutionary competition, at least as a coordination rule for games with strict (weakly) Pareto optimal Nash equilibria. Yet it does not show that hypothetical bargaining is the decision rule which outcompeted other decision rules with the same or better resistance to invasions.

According to the proponents of the hypothetical bargaining theory, hypothetical bargaining is a mode of reasoning which allows people to coordinate their actions (Misyak and Chater 2014, Misyak et al. 2014). More specifically, it is a set of inference rules which allows the decision-makers to identify the bargaining solution of a game – an outcome which has a specific set of properties. The set of hypothetical bargaining solutions of a game is supposed to be smaller than the set of all the possible outcomes. Assuming that the set of feasible agreements is the set of Nash equilibria, the set of hypothetical bargaining solutions of a game will be a subset of the Nash equilibria. In some games, that subset will be considerably smaller than the set of Nash equilibria. Therefore, by choosing to play their part in the attainment of a bargaining solution of a game, hypothetical bargainers either resolve the coordination problem completely or, in games where the bargaining solution is not unique, increase the probability of successful coordination and, consequently, their expected payoff. Even in games where players have conflicting preferences over the Nash equilibria of the game, playing a part in the attainment of a hypothetical bargaining solution may be the best option for everyone solely due to higher probability of coordination success.

In addition, hypothetical bargaining is a mode of reasoning which incorporates certain benefit distribution considerations. Recall that hypothetical bargainers are assumed to search for an efficient solution of a game which each self-interested player would be motivated to accept. Like the team reasoning theory, hypothetical bargaining theory suggests that players seek for mutually advantageous solutions of games: Hypothetical bargainers always choose a (weakly) Pareto optimal outcome, provided that there is a (weakly) Pareto optimal outcome which satisfies the feasibility criterion. According to Bacharach (2006), people's ability to coordinate their actions in the at-

tainment of efficient outcomes can be explained as an outcome of the group selection process. A group of individuals whose members have the ability to coordinate their actions in the attainment of efficient outcomes should have a higher average fitness than the group whose members have no such ability (provided, of course, that individuals belonging to a group whose members have the efficient coordination ability interact with each other sufficiently frequently). Thus, a group of individuals with efficient coordination ability would tend to grow faster than other groups (for details, see Bacharach 2006).

For the aforementioned reasons, it is worthwhile to explore the possibility that hypothetical bargainers were able to spread in the population because hypothetical bargaining was a reasoning procedure which allowed the hypothetical bargainers to resolve interdependent decision problems efficiently, and thus gave them an evolutionary advantage against individuals who used other types of decision-making procedures.

It seems plausible that hypothetical bargaining started as a type of *familiarity heuristic* – a relatively primitive decision rule based on decision-maker’s recognition of certain similarities between different decision problems and/or decision contexts. More specifically, a decision-maker using a familiarity heuristic recognizes that a new or less frequently encountered decision problem is, in some respect, similar to another recurrent decision problem, the solution of which is already known to the decision-maker. Decision-maker’s recognition of the similarities between the new or less frequently encountered decision-problem and the recurrent decision problem with the known solution prompts him or her to choose a solution of a new or less frequently encountered decision problem which is, in some respect, similar to the already known solution of the recurrent decision-problem (for extensive discussion, see Ashcraft 2006, Tversky and Kahneman 1974).

There seems to be several reasons to believe that prehistoric decision-makers could have developed the ability to recognize the structural similarities between resource division problems and other types of decision problems.

Various types of conflicts over divisions of limited resources is one of the most common problems faced by every human population, irrespective of its size or level of cultural development. In fact, various animal species, especially those capable of engaging in collective actions or exhibiting basic level of social organization, also face this type of problem. It seems reasonable to believe that prehistoric human populations had to deal with the resource division problem very frequently. Some kind of cognitive response to this type of problem had to emerge relatively early, at least in those populations which managed to survive for an extended period of time. In other words,

individuals had to develop the ability to recognize this type of problem, and respond to it by taking a certain course of action. Given the frequency with which individuals had to face this type of problem, it is possible that individuals developed a habitual response to it.

This assumption is indirectly supported by the empirical fact that even animal species of lesser intelligence than monkeys are capable of following certain food allocation algorithms. For example, it has been observed that packs of wolves follow a hierarchy-based procedure where the pack leader (alpha animal) gets access to food first, and is then followed by beta animals, gamma animals, and so on (Mech 1999). The prehistoric human populations probably followed much more sophisticated resource division rules than those followed by wolves or monkeys.

Assuming that a habitual response to resource division problems was one of the first to develop, it seems reasonable to believe that some individuals used the same response to other, possibly newly encountered, types of decision problems, most likely to problems for which a habitual response was not available at that time and/or which shared some structural similarities with resource division problems.

An instinct of survival should drive the individual to use any means necessary to secure the resources which are essential for survival, even if those means involve an infliction of a physical harm to other individuals and/or the risk of getting physically harmed by others. It stands to reason to believe that an individual receiving a lesser than necessary share of resources should be more likely to engage in aggressive actions towards other individuals, even at the expense of his or her own safety. This is likely to occur in situations where an individual would interpret the actions of other individuals as a takeover of the basic resources necessary for individual's survival.

In the absence of any sophisticated measures of the sizes of resource shares, individual's response in the resource division games was likely based on crude comparisons of the shares of resources received by each interacting individual. It seems reasonable to believe that an individual was more likely to engage in aggressive behaviour in situations where s/he recognized a disadvantageous inequality between the shares of resources received by the interacting parties. In other words, an individual was more likely to engage in aggressive behaviour if s/he recognized that his or her share of the basic resources was smaller than the shares of resources gained by others.

In the absence of an external arbitrator who can keep the actions of conflicting individuals in check, the conflicts over division of resources may have considerable destructive power, especially if physical violence is involved between individuals of roughly similar physical abilities. Frequent physical

conflicts pose a risk for individual's and even population's survival. Due to high frequency with which the prehistoric societies had to face the resource division problem, and due to a looming threat of each conflict over resources escalating to destructive levels, it seems likely that individuals developed an impulse to respond to this problem in a way which reduced the likelihood of destructive conflicts. In order to resolve such conflicts in a way which minimized the likelihood of violent conflicts, individuals had to develop the ability to take other individual's reactions to their own actions into account. More specifically, individuals had to develop the ability to recognize the relationship between their actions and their opponents' reactions, and to avoid the types of actions which triggered the dangerous reactions of other individuals. Since each individual's violent reactions were likely associated with unequal divisions of resources, individuals had developed an impulse to avoid such unequal divisions, at least between individuals of roughly similar physical and mental abilities, possibly of roughly similar standing in the social hierarchy. The benefit distribution considerations, which underlie bargaining-based approach to decision-making, may be based on fundamental self-preservation considerations, not on pro-social motivations.

The preceding arguments are, by no means, decisive. However, it is difficult to deny an intuition that our understanding of what is the appropriate solution of an interdependent decision problem is influenced by some deeply ingrained collective acceptability considerations. Resource allocation problem was, almost certainly, one of the first interdependent decision problems that prehistoric societies faced frequently. Therefore, a hypothesis that individuals' reasoning about interdependent decision problems in general was shaped by their repeated engagement in resource division games seems plausible.

Even if hypothetical bargaining offers a plausible story of how some individuals developed the ability to resolve interdependent decision problems, there is no obvious reason to believe that hypothetical bargainers out-competed other types of decision-makers and became the dominant type of decision-maker. The success of hypothetical bargaining in evolutionary competition with other decision rules would depend on the initial population conditions – the types of decision-makers that hypothetical bargainers had to compete with. Hypothetical bargainers can be convincingly shown to be advantageous when competing with some extremely primitive forms of behaviour, but results are far from conclusive when more sophisticated decision-makers are assumed to be present in the population.

For example, suppose that population plays a three strategy coordination game depicted in Figure 4.3. This game has three pure strategy Nash equilib-

	$s1$	$s2$	$s3$
$s1$	7, 3	0, 0	0, 0
$s2$	0, 0	5, 5	0, 0
$s3$	0, 0	0, 0	3, 7

Figure 4.3: Coordination game

ria $(s1, s1)$, $(s2, s2)$ and $(s3, s3)$ and four Nash equilibria in mixed strategies³³. It will be assumed that each player's expected payoff represents his or her fitness – the expected number of offspring or imitators³⁴.

Let us assume that a population consists of unsophisticated decision-makers who choose one of the pure strategies at random (that is, behave like the level 0 players in the cognitive hierarchy theory). The incumbent decision rule thus prescribes strategy $x = \left(\frac{1}{3}s1, \frac{1}{3}s2, \frac{1}{3}s3\right)$ to every interacting individual. Each individual's payoff from using the incumbent decision rule is thus

$$u(x, x) = \left(\left(\frac{1}{9} \times 7\right) + \left(\frac{1}{9} \times 5\right) + \left(\frac{1}{9} \times 3\right)\right) = \frac{15}{9} = \frac{5}{3}. \quad (4.22)$$

Suppose that a small share $\varepsilon > 0$ of hypothetical bargainers appears in the population. Each hypothetical bargainer identifies the Nash equilibrium $(s2, s2)$ as the solution of this game and thus always plays strategy $s2$. Assuming that players' matching is random, hypothetical bargainer's expected payoff from using hypothetical bargaining decision rule in a new population state is

$$u(s2, \varepsilon s2 + (1 - \varepsilon)x) = 5\varepsilon + \frac{5}{3}(1 - \varepsilon). \quad (4.23)$$

The unsophisticated decision-maker's payoff from using an incumbent decision rule is

$$u(x, \varepsilon s2 + (1 - \varepsilon)x) = \frac{5}{3}\varepsilon + \frac{5}{3}(1 - \varepsilon). \quad (4.24)$$

Notice that the equality $\frac{5}{3}\varepsilon + \frac{5}{3}(1 - \varepsilon) = \frac{5}{3}$ holds for any $\varepsilon \in (0, 1)$.

Hypothetical bargainers will spread in the population if the following condition is satisfied:

$$u(s2, \varepsilon s2 + (1 - \varepsilon)x) > u(x, \varepsilon s2 + (1 - \varepsilon)x). \quad (4.25)$$

³³The two mixed strategy Nash equilibria of this game are: $\left(\frac{5}{8}s1, \frac{3}{8}s2; \frac{5}{12}s1, \frac{7}{12}s2\right)$, $\left(\frac{7}{10}s1, \frac{3}{10}s3; \frac{3}{10}s1, \frac{7}{10}s3\right)$, $\left(\frac{7}{12}s2, \frac{5}{12}s3; \frac{3}{8}s2, \frac{5}{8}s3\right)$, $\left(\frac{35}{71}s1, \frac{21}{71}s2, \frac{15}{71}s3; \frac{15}{71}s1, \frac{21}{71}s2, \frac{35}{71}s3\right)$

³⁴This model can be interpreted as a biological reproduction model, as well as a model of cultural imitation.

It is easy to check that $5\varepsilon + \frac{5}{3}(1 - \varepsilon) > \frac{5}{3}$ for any $\varepsilon > 0$. Hypothetical bargainers could therefore be expected to spread in this population of unsophisticated decision-makers.

In the aforementioned example, hypothetical bargainers do not need to have a perfectly developed ability to recognize the outcome $(s2, s2)$ as the bargaining solution. Suppose that each hypothetical bargainer's ability to recognize outcome $(s2, s2)$ is imperfect, but s/he manages to do that with some positive probability $\mu > 0$ which is the same for every player of this type. With probability $(1 - \mu)$, each hypothetical bargainer chooses one of the pure strategies at random as other unsophisticated decision-makers.

Suppose that $\varepsilon > 0$ is the share of hypothetical bargainers, while $(1 - \varepsilon)$ is the share of unsophisticated decision-makers. For each hypothetical bargainer, the probability of being matched with another hypothetical bargainer is ε . If two bargainers are drawn to play the game, they will both choose strategy $s2$ with probability μ^2 . With probability $\mu(1 - \mu)$, one of them will play strategy $s2$ while the other one will choose one of the pure strategies at random. Finally, with probability $(1 - \mu)^2$ both bargainers will choose their strategies randomly.

The expected payoff of each hypothetical bargainer from his or her interaction with the other hypothetical bargainer is

$$u_b = \left(5\mu^2 + \frac{5}{3}\mu(1 - \mu) + \frac{5}{3}(1 - \mu)^2 \right). \quad (4.26)$$

Recall that each unsophisticated decision-maker's expected payoff is $\frac{5}{3}$ for any $\varepsilon \in (0, 1)$. Hypothetical bargainer's expected payoff from playing against the unsophisticated decision-maker is $\frac{5}{3}$ for any $\varepsilon \in (0, 1)$. It follows that hypothetical bargainer's strategy yields the same expected payoff when matched with the strategy of an unsophisticated decision-maker as unsophisticated decision-maker's strategy when matched with itself. It means that hypothetical bargainers will be able to spread in the population if hypothetical bargainer's strategy yields a higher expected payoff when matched with itself than the strategy of an unsophisticated decision-maker when matched with the hypothetical bargainer's strategy:

$$5\mu^2 + \frac{5}{3}\mu(1 - \mu) + \frac{5}{3}(1 - \mu)^2 > \frac{5}{3}. \quad (4.27)$$

A relatively simple computation reveals that the inequality will hold for any $\mu > \frac{1}{3}$ ³⁵. In other words, a hypothetical bargainer must only be able to iden-

³⁵It is easy to check that equality $5\mu^2 + \frac{5}{3}\mu(1 - \mu) + \frac{5}{3}(1 - \mu)^2 = \frac{5}{3}$ holds when $\mu = 0$ and $\mu = \frac{1}{3}$. The output of function $f(\mu) = \left[\left(5\mu^2 + \frac{5}{3}\mu(1 - \mu) + \frac{5}{3}(1 - \mu)^2 \right) - \frac{5}{3} \right]$ is positive for any $\mu > \frac{1}{3}$. The probability of each bargainer getting the impulse to choose strategy $s2$ in order to be evolutionarily successful must be higher than $\frac{1}{3}$.

tify the bargaining solution with a probability of higher than 1/3 in order for hypothetical bargainers to outcompete the unsophisticated decision-makers.

Assuming that hypothetical bargainer's ability to identify the bargaining solution is improving over time (meaning that μ is increasing over time), their advantage over unsophisticated players will increase as well, meaning that hypothetical bargainers will spread in the population more rapidly. If the population reaches a state where everyone is a perfect hypothetical bargainer (that is, a hypothetical bargainer, such that $\mu = 1$), individuals will play a strict Nash equilibrium (s_2, s_2) in every interaction, and such a population state is evolutionarily stable.

The preceding example offers a scenario where bargainers could be expected to spread in the population. In that scenario, a population state where everyone is using an incumbent decision rule is not an evolutionarily stable state: A strategy profile in which each decision-maker plays one of the pure strategies with probability 1/3 is not a Nash equilibrium (for an extensive technical discussion of the relationship between the concept of Nash equilibrium and the concept of evolutionary stability, see Weibull 1995). The model could be criticized for being based on 'convenient' assumptions concerning the behaviour of unsophisticated decision-makers.

Although there is no obvious reason to believe that a population of unsophisticated players must necessarily be in an evolutionarily stable state when hypothetical bargainers appear, one could argue that such a scenario is less likely than the one where unsophisticated players are in some evolutionarily stable state: A population of unsophisticated players would, even by a sequence of accidental deviations from the standard behaviour (in the aforementioned example, from the choice of each pure strategy with probability 1/3), end up playing a Nash equilibrium of the game, since evolutionary successful deviants would take over the population³⁶.

There are reasons to believe that hypothetical bargainers would be much less successful in a population of more sophisticated decision-makers. In other words, other types of decision-makers could be following simple yet evolutionary competitive decision-making rules. In addition, the population of more sophisticated decision-makers could already be in an evolutionarily stable state, thus making the invasion of hypothetical bargainers far less likely.

In such cases, the evolutionary success of hypothetical bargainers would depend on a number of factors, such as their share in the initial population state, their ability to successfully identify the bargaining solutions of the game, the structure of the game, and the types of decision-making rules

³⁶For an extensive discussion of this view, see Binmore 2005.

followed by other individuals in the population.

For example, suppose that a population is playing a three strategy mixed-motive game depicted in Figure 4.4. Notice that this game is structurally

	s	r	e
s	5, 5	0, 2	1, 4
r	2, 0	2, 2	2, 0
e	4, 1	0, 2	0, 0

Figure 4.4: Mixed motive game

similar to the Stag Hunt game: It has two pure strategy Nash equilibria (s, s) and (r, r) , and a mixed strategy Nash equilibrium $(\frac{2}{5}s, \frac{3}{5}r; \frac{2}{5}s, \frac{3}{5}r)$.

Suppose that some individuals in the population are hypothetical bargainers – decision-makers who always get an impulse to choose strategy s . Another type of individual is a ‘survivalist’ – a type of decision-maker who gets an impulse to choose an action which guarantees a positive payoff, irrespective of what the other player does. An assumption that some individuals in the population have this impulse seems rather plausible, since such a strategy is likely to ensure survival in an environment where the resources necessary for survival are scarce. In game theoretic terms, a survivalist always plays a *maximin* strategy r , which is also the risk-dominant strategy of this game.

Finally, the third type of individual is an ‘exploiter’ – a type of decision-maker who initially mimics cooperative behaviour, yet later gets an impulse to force the players engaged in genuine cooperative actions to invest most of their efforts into getting the resources for the exploiter rather than for themselves. This strategy will be assumed to be useless against survivalists and other exploiters.

It will be assumed that every type of player always gets the same impulse and follows it perfectly: The bargainers always choose strategy s , the survivalists always choose strategy r , and the exploiters always choose strategy e .

From this set of assumptions, it follows that a population state where every individual is a hypothetical bargainer will be playing a strict Pareto efficient Nash equilibrium (s, s) , while a population state where every individual is a survivalist will be playing a strict risk-dominant Nash equilibrium (r, r) . Both population states are evolutionarily stable.

A global dynamic picture under the replicator dynamics is shown below in Figure 4.5(a). A global dynamic picture under the best response dynamics

is shown in Figure 4.5(b)³⁷.




Figure 4.5: Mixed motive game under (a) replicator dynamics and (b) best response dynamics

The three vertices of the simplex represent three extreme population states: One in which everyone plays s , one in which everyone plays r , and one in which everyone plays e . The dots in both diagrams represent the stationary states, or rest points. The black dots represent the evolutionarily stable states, while the white dots represent the unstable stationary points. The diagram shows that this game has two evolutionarily stable states: One in which everyone plays strategy s and one in which everyone plays strategy r . The white dot on the left edge of the simplex represents a mixed stationary population state where $\frac{2}{5}$ of the population plays strategy s , and $\frac{3}{5}$ of the population plays strategy r . This stationary state is not evolutionarily stable, meaning that a small perturbation of shares of individuals playing these strategies would move the population to another stationary state. The arrows indicate the trajectories of the evolutionary dynamics.

Both dynamics show that the population may be in one of the two evolutionarily stable states: In one state everyone is a hypothetical bargainer, while in the other one everyone is a survivalist. The white dot on the left edge of the simplex represents a stationary mixed population state where $\frac{2}{5}$ of the population are hypothetical bargainers and $\frac{3}{5}$ are survivalists. The

³⁷All the simulations and diagrams have been produced with *Dynamo* package for Wolfram Mathematica developed by Sandholm, Documaci and Franchetti (2012).

dynamics shows that exploiters would not survive in the competition against hypothetical bargainers and survivalists, and that a population state where everyone is an exploiter is not evolutionarily stable. Exploiters should thus be driven out of the population from any initial population state. Both dynamics show that if the share of hypothetical bargainers in the initial population state were larger than $\frac{2}{5}$, this strategy would outcompete the survivalists and take over the population. This also means that if every individual in the population were a survivalist (meaning that the population was in the evolutionarily stable state where everyone plays r), the hypothetical bargainers could invade the population, but only if the share of mutants were larger than $\frac{2}{5}$ of the total population. An event of such a large number of survivalists simultaneously mutating into hypothetical bargainers seems to be fairly unlikely. It is easy to check that the proportion of hypothetical bargainers needed for the strategy to take over the population decreases with increasing payoff gains associated with cooperative behaviour³⁸.

The evolutionary success of hypothetical bargaining also depends on how individuals update their strategies. For example, in a population of myopic best responders in which exploiters are extremely efficient, hypothetical bargainers could neither compete against the survivalists, nor could they invade such a population. For example, suppose that the exploiter mimics the behaviour of a cooperative individual, yet later takes all the cooperative product and leaves a genuine cooperator without any resources. Such a strategic interaction may look like the game depicted in Figure 4.6. This game also

	s	r	e
s	5, 5	0, 2	0, 5
r	2, 0	2, 2	2, 0
e	5, 0	0, 2	0, 0

Figure 4.6: Mixed motive game with efficient exploiters

has two pure strategy Nash equilibria (s, s) and (r, r) , and six mixed strat-

³⁸For example, suppose that, *ceteris paribus*, the payoffs associated with the Nash equilibrium (s, s) increases from 5 to 10 (say, due to increased production efficiency associated with new technology). The mixed strategy Nash equilibrium of the new game would become $(\frac{1}{5}, \frac{4}{5}; \frac{1}{5}, \frac{4}{5})$. Given this new game, the hypothetical bargainers would take over the population if their share in the population were strictly larger than $\frac{1}{5}$ of the total population.

egy Nash equilibria³⁹. The global dynamic picture under the best response dynamics is shown in Figure 4.7(a).




Figure 4.7: Mixed motive game with efficient exploiters under (a) best response dynamics and (b) replicator dynamics.

In this case there is only one evolutionary stable population state – a state where everyone is a survivalist playing strategy r . Notice that hypothetical bargainers would not drive the survivalists from the population if the destructive exploiters were present, even in cases where the share of bargainers in the initial population state were large. The dynamics shows that, due to exploitative strategy being advantageous against the cooperative strategy of hypothetical bargainers, the hypothetical bargainers would swiftly change their strategy into exploitative strategy and the population would shift to a state with more exploiters (the arrows on the upper part of the simplex indicate the direction of this process). However, this state is not stationary, which means that the population would not stay in that state for long. In the long run, the survivalists would take over the population. To see why this would happen, notice that hypothetical bargainers would get no payoff from interacting with exploiters and therefore, when given a chance, they would change their strategy into exploitative strategy. The share of hypothetical bargainers would drop. The share of exploiters, however, drops

³⁹The six mixed strategy Nash equilibria of this game are: $(s; \frac{2}{5}s, \frac{3}{5}e)$, $(\frac{2}{5}s, \frac{3}{5}r; \frac{2}{5}s, \frac{3}{5}r)$, $(\frac{2}{5}s, \frac{3}{5}r; \frac{2}{5}s, \frac{3}{5}e)$, $(\frac{2}{5}s, \frac{3}{5}e; s)$, $(\frac{2}{5}s, \frac{3}{5}e; \frac{2}{5}s, \frac{3}{5}e)$, $(\frac{2}{5}s, \frac{3}{5}e; \frac{2}{5}s, \frac{3}{5}e)$.

due to their strategy being ineffective against the strategy of survivalists. The exploiters are thus motivated to adopt the strategy of survivalists. In the end, the survivalists would dominate the population.

If the players were not able to update their strategies using such a relatively sophisticated rule, the hypothetical bargainers could be much more successful. For example, consider the global dynamic picture of the same game under the replicator dynamics depicted in Figure 4.7(b). The dynamics shows that multiple mixed population states in which hypothetical bargainers are present are evolutionarily stable (notice the mixed evolutionarily stable states in which survivalists are absent). In addition, a population state in which everyone is a hypothetical bargainer is an evolutionarily stable state. The analysis reveals that hypothetical bargainers could spread from any initial population state in which more than $\frac{2}{5}$ of the total population were hypothetical bargainers. Hypothetical bargainers could also invade a population in which everyone were a survivalist, yet the number of mutants would have to be larger than $\frac{2}{5}$ of the population. This example suggests that a higher sophistication of individuals may not work in favor of hypothetical bargaining.

In conclusion, evolutionary game theory offers only a limited support to the idea that hypothetical bargaining evolved as an evolutionary response to coordination problems. Interpreted as a simple decision rule, hypothetical bargaining can be shown to successfully spread in populations where other types of individuals are using extremely primitive responses to coordination problems (e.g. choose their strategies at random). In evolutionary games where hypothetical bargainers play a strict Nash equilibrium, a population state in which everyone is a hypothetical bargainer is evolutionarily stable. This means that, in principle, hypothetical bargaining could have emerged as an evolutionary response to certain types of interdependent decision problems and resisted the invasions of other decision rules.

However, in populations where other types of individuals are using more sophisticated decision rules, hypothetical bargaining can spread only if their share in the initial population state is sufficiently large to begin with. This also means that, in many cases, hypothetical bargainers could only invade the population if a sufficiently large number of individuals were to adopt this decision rule simultaneously. The question of why a sufficiently large number of individuals would adopt this decision rule simultaneously when playing a particular game cannot be answered with the purely formal analytic tools provided by evolutionary game theory. Several speculative responses could be provided. One possibility is that hypothetical bargaining emerged among unsophisticated decision-makers who were playing games in which

hypothetical bargaining could have easily outcompeted the unsophisticated responses. This speculation, however, does not explain why hypothetical bargaining rather than some other possible decision rule would have invaded such a population. Another possible response, suggested by Bacharach 2006, is that players capable of achieving efficient coordination were interacting more frequently with each other than with other types of players, thus allowing them to gain advantage over other types of individuals. In Bacharach's evolutionary story, populations are assumed to be divided into groups of individuals. Individuals belonging to the same group are assumed to interact more frequently with each other than with individuals from other groups. A group containing more individuals capable of achieving efficient coordination would have a higher average fitness than a group of individuals lacking this ability. Consequently, a group containing more efficient coordinators would grow faster, thus outcompeting other groups (for extensive discussion, see Bacharach 2006).

This story has some conceptual credibility: Skyrms (1996) has shown that a correlation assumption (i.e. an assumption that individuals of the same type interact more frequently with each other than with other types of individuals) may indeed eliminate certain evolutionarily stable mixed population states, and work in favor of certain types of strategies by reducing the number of players which would have to be using that strategy in the initial population state in order for it to spread in the population (for details, see Skyrms 1996). Without empirical support, however, such explanations remain highly speculative, since they postulate certain assumptions about the prehistoric population which cannot be evaluated empirically (for a critical discussion of Skyrms's correlation assumption, see D'Arms 1996).

It is important to note that the aforementioned analysis was based on the assumption that decision rule's evolutionary success in a particular game depends purely on the success of its prescribed strategy. It is possible that such a simplistic behavioural interpretation of evolutionary selection of decision rules is inadequate to represent the evolutionary selection of reasoning modes. However, without an adequate model of mental bargaining process, a more sophisticated interpretation of hypothetical bargaining is currently not feasible.

4.5 The Non-Uniqueness Problem

Even if hypothetical bargaining could be developed into a conceptually sound model of strategic reasoning, there are serious reasons to believe that it would not be a single generalizable model from which the theoretical predictions of

players' choices in any non-cooperative game could be derived. The reason lies in the nature of non-cooperative game itself. In standard bargaining problems, two players have to decide on how to split a divisible good. Each player's utility function represents his or her preferences over the feasible allocations of the good, and the bargaining solution is the unique distribution of the good which satisfies a specific set of desirable properties. In non-cooperative games, however, players' utility functions may represent any kind of personal motivations, and so there may be multiple outcomes in a game having the same set of desirable formal properties. Hypothetical bargainer's choice among the multiple outcomes which s/he identifies as having the set of desirable properties cannot be derived from the formal model of hypothetical bargaining. Hypothetical bargainers' ability to coordinate their actions in such situations would thus likely be based on conformity 'to complex and sometimes arbitrary conventions that could not be reconstructed by abstract rational analysis' (Sugden 2015: 156). This seems to be acknowledged by Misyak et al. 2014 who argue that successful coordination will depend on players' common knowledge of previous hypothetical bargains, since 'common knowledge of the precedent marks that coordination solution as "special" and thus acts as a possible tiebreaker for choosing between future solutions' (Misyak et al. 2014: 516).

For example, suppose that two individuals are participating in an experiment where they are asked to pick a strategy in the extended Hi-Lo game which is presented to them in the form depicted in Figure 4.8:

	$r1$	$r2$	$r3$
$r1$	10 * 10	0, 0	0, 0
$r2$	0, 0	10, 10	0, 0
$r3$	0, 0	0, 0	5°5

Figure 4.8: Extended Hi-Lo game

For simplicity, suppose that both individuals are hypothetical bargainers and this fact is common knowledge. Hypothetical bargainers may believe that in open negotiations they would easily agree to attain either outcome $(r1, r1)$ or outcome $(r2, r2)$: In terms of formal properties, both outcomes are identical. Without communication, however, the players could not send each other any signals which would allow them to coordinate their actions in the attainment of one of the two formally identical bargaining solutions. This means that hypothetical bargainers would face a coordination problem.

In some games, hypothetical bargainers could resolve this problem by taking into consideration the perceived coordination success rate, and then consider playing *ex ante* Pareto dominated outcome if, given the coordination success rate, the *ex ante* Pareto dominated option yields a higher expected payoff for every hypothetical bargainer⁴⁰. In this game, however, this approach would not resolve the problem: From playing strategies $r1$ and $r2$ at random (that is, with probability $1/2$), each player would get the same expected payoff of 5 as s/he would get from playing strategy $r3$.

Hypothetical bargainers could notice that the Pareto optimal outcome $(r1, r1)$ is marked with a star, while the other one is not. The players could thus attempt to use this star as a coordination aid. Hypothetical bargainers could also use the absence of a star to coordinate their actions on the Pareto optimal outcome $(r2, r2)$. Finally, the players could recognize that the outcome $(r3, r3)$ is marked with a diamond, and use it (or its absence) as a coordination aid as well. Each of the arbitrary attributes which has nothing to do with the payoff structure of the game (or its absence) could potentially be used as a coordination aid. Hypothetical bargainer's choice of the coordination aid will depend on his or her beliefs about which one of the aids is most likely to be recognized, considered, and adopted by another hypothetical bargainer. These beliefs may in turn depend on hypothetical bargainers' cultural backgrounds, the prevalent social conventions in their societies, and other factors unrelated to the formal game theoretic properties of the game itself. Thus, a formal model of hypothetical bargaining which could predict players' *final* choices in such context-sensitive decision problems may not be possible at all.

4.6 Conclusion

Hypothetical bargaining theory suggests that people choose their strategies in non-cooperative games on the basis of what they believe they would agree to play if they could openly bargain. The hypothetical bargaining should be viewed as a goal-directed mode of reasoning – a set of inference rules which allows the players to identify the outcomes of a game with a specific set of desirable properties.

The proponents of the theory suggest that hypothetical bargaining is compatible with the principles of rational choice theory. For this interpretation of hypothetical bargaining to be credible, it must be possible to show that hypothetical bargainer's decision to play a part in the attainment of the

⁴⁰In the context of team reasoning theory, this idea has been discussed by Bardsley et al. 2010 and Faillo et al. 2016.

outcome identified as the bargaining solution of a game is compatible with the orthodox choice consistency principles.

In this chapter I have argued that there are several reasons of why the rational choice interpretation of hypothetical bargaining is conceptually problematic. The first reason is hypothetical bargaining theory's vulnerability to choice rationalization problem: Common knowledge of the fact that each player is a hypothetical bargainer and has identified a certain outcome as the hypothetical bargaining solution of the game does not give the hypothetical bargainer a rational reason to play his or her part in the attainment of that outcome. Hypothetical bargainer's decision to choose a part in the attainment of the outcome is rational only if s/he expects the other players to do that as well. The theory does not offer an explanation of why each hypothetical bargainer should be expected to hold such a belief, and so does not provide an explanation of why a hypothetical bargainer should be motivated to actually play his or her part in realizing the outcome which s/he recognizes as the hypothetical bargaining solution of a game.

I have argued that this difficulty could, in principle, be resolved if hypothetical bargaining were to be interpreted as a belief-formation algorithm rather than as a choice algorithm, and suggested that lexicographic belief systems used in epistemic rationalizability models could, in principle, serve as an imperfect yet useful starting point for the development of an adequate formal representation of such a belief-formation algorithm. However, I have also indicated that such an interpretation of hypothetical bargaining involves a non-trivial departure from the epistemic principles of the orthodox game theory.

The second reason is that hypothetical bargaining, if interpreted as a rational choice explanation of social coordination, can only account for the actions of individuals who express a common belief that the share of hypothetical bargainers in the population is sufficiently large. I have argued that such a belief can only form in a population where a sufficiently large number of individuals are acting as hypothetical bargainers. If interpreted as a rational choice explanation of social coordination, the theory seems to fall into a vicious cycle: Hypothetical bargainers can only be motivated to act in a population state that cannot be reached without a sufficient number of players acting as hypothetical bargainers. I have argued that the most plausible explanation of how such a population state could have emerged in the first place is that hypothetical bargaining first emerged as a familiarity heuristic which was able to spread in the population due to its fitness-enhancing properties.

However, evolutionary game theoretic models considered in this chapter

only show that individuals using such a familiarity heuristic could have been more successful than individuals using certain extremely primitive decision rules. I have argued that the question of whether basic evolutionary game theory models are adequate for representing the actual evolution of decision-making rules cannot be answered with the tools of the evolutionary game theory itself, and so the evolutionary game theoretic explanation suggested in this chapter is valid insofar as it shows that an intuition that people's choices in non-cooperative games may be driven by an evolved sensitivity to benefit distribution considerations has some conceptual credibility.

Finally, decision-makers' ability to coordinate their actions in games with multiple hypothetical bargaining solutions would depend on factors that are not related to payoff structures of games alone, such as shared cultural norms and conventions. Therefore, even if hypothetical bargaining were developed into a conceptually sound model of strategic reasoning, a development of a single generalizable model of hypothetical bargainers' *final* choices in non-cooperative games would not necessarily be possible.

Chapter 5

Conclusion

Orthodox game theoretic analysis of non-cooperative games is based on best-response reasoning. It is often criticized for producing multiple solutions, even in games which seem to have intuitively obvious unique solutions and create no coordination problems for the real-world decision-makers. This prompted the emergence of multiple theories which purport to explain how people resolve games that, from the perspective of orthodox game theory, have multiple rational solutions. Two of the more recent theories – the team reasoning theory and the hypothetical bargaining theory – suggest that people resolve non-cooperative games by following a reasoning procedure which allows them to identify the mutually advantageous solutions. A number of properties have been suggested in the literature for a mutually advantageous solution to satisfy: Pareto efficiency, feasibility, successful coordination of the interacting decision-makers' actions, and equitable distribution of individuals' personal payoff gains. Yet so far very few formal characterizations of the concept of mutual advantage which could be incorporated into the formal game theoretic analysis have been proposed.

In this thesis, I have suggested two possible formal characterizations of mutual advantage which could be derived from the aforementioned theories: The notion of mutual advantage as the maximization of the minimum level of individual advantage among the interacting players, developed in collaboration with Karpus (Karpus and Radzvilas 2016), which is broadly in line with the notion of mutual advantage suggested in Sugden's (2011, 2015) version of the team reasoning theory, and the benefit-equilibrating hypothetical bargaining solution, which I have suggested as a possible formal characterization of outcomes which hypothetical bargainers would identify as mutually advantageous and agreeable solutions of games. I have discussed their formal properties and theoretical predictions in a number of experimentally relevant games.

The solution concepts suggested in this thesis share certain conceptual similarities: Both of them are based on the principle that decision-makers will never identify a feasible solution as mutually advantageous if there is another feasible solution which is strictly better for every interacting player. Weak Pareto efficiency seems to be one of the natural properties that any mutually advantageous solution of a game should satisfy. The differences between the properties of the suggested solution concepts reflect the differences between the core assumptions of the theories from which they are derived.

According to the team reasoning theory, certain structural and/or conceptual features of games may trigger a shift in decision-maker's mode of reasoning from individualistic reasoning to reasoning as a member of a team. This shift of reasoning involves a transformation of agency: A team-reasoning decision-maker identifies himself or herself with a group of individuals who act together in the attainment of some common goal. More specifically, team-reasoning decision-maker identifies the attainment of his or her goal in a game with team's success in the attainment of its goal. This creates a personal motivation for a team-reasoning decision-maker to play his or her part in the attainment of an outcome that s/he recognizes as the team's goal. According to team reasoning theory, team's goal may not be the maximum advancement of the personal interests of every member of a team, and so team reasoning decision-maker's choice to play a part in the attainment of the team's goal may not lead to the maximum advancement of his or her personal interests. In other words, in order to play a part in the attainment of the team's goal, a decision-maker who reasons as a member of a team may choose a strategy which, given the expected actions of other team-reasoning decision-makers, does not lead to the maximum advancement of decision-maker's own personal interests.

Because of this, the theory of team reasoning predicts an out-of-equilibrium play in certain experimentally relevant games, such as the Prisoner's Dilemma game. It also offers an explanation of why people cooperate in social dilemmas – interdependent decision problems where individual incentives and social optimality diverge.

Yet according to the principles of orthodox game theory, team-reasoning decision-makers' actions could be interpreted as being driven by a structure of incentives which is different from the structure of incentives which is represented by the payoff structure of the original game. The theory of team reasoning thus allows for modelling of team-reasoning decision-maker's incentive to play a part in the attainment of the team's goal as being independent from decision-maker's personal incentives that motivate his or her actions before a shift from individualistic best-response reasoning mode to

team mode of reasoning occurs.

This feature of the team reasoning theory is also reflected in the notion of mutual advantage as the maximization of the minimum extent of individual advantage among the interacting players. In a considerable number of games, there may be multiple outcomes with this property, yet each outcome may be associated with a different allocation of players' individual advantage gains. A team-reasoning decision-maker would identify every outcome with the aforementioned property as leading to the attainment of the team's goal, irrespective of how individual advantage gains were distributed among the interacting players. This implies that a team-reasoning decision-maker who aims to maximally advance the interests of a team would be indifferent between playing a part in the attainment of a team optimal outcome associated with a lower level of his or her own individual advantage and playing a part in the attainment of a team optimal outcome associated with a higher personal advantage gain.

Unlike the team-reasoning decision-makers, hypothetical bargainers are assumed to be self-oriented decision-makers – individuals who, like the best-response reasoners, aim to maximize their individual advantage as much as possible. A hypothetical bargainer expects the other hypothetical bargainers to deviate from any outcome which creates them a personal incentive to do so. This expectation restricts the set of outcomes that hypothetical bargainers deem implementable via their joint actions. A hypothetical bargainer only cares about the personal interests of the other players insofar as their actions may promote or hinder the advancement of his or her personal interests, and expects the other hypothetical bargainers to have similar motivations as s/he does. That is, s/he expects every other hypothetical bargainer to prefer a more personally advantageous feasible outcome over any less personally advantageous feasible outcome. A hypothetical bargaining solution, such as the benefit-equilibrating solution suggested in this thesis, represents hypothetical bargainer's expectation of how self-oriented decision-makers would agree to distribute their maximum attainable individual advantage losses in order to reach an agreement to implement a weakly Pareto optimal feasible outcome rather to end up with no agreement on how to play the game at all. Hypothetical bargainers may use this common expectation as a resolution of a game with multiple Nash equilibria.

Because of the assumption that hypothetical bargainers are individualistic decision-makers, the hypothetical bargaining theory makes a less radical departure from individualistic non-cooperative game theory than the team reasoning theory. It offers an individualistic explanation of how people resolve non-cooperative games by identifying mutually advantageous and

agreeable outcomes. In addition, the benefit-equilibrating solution concept suggested in this thesis seems to better account for people's choices in certain experimentally relevant coordination games than the theory of team reasoning based on the notion of mutual advantage suggested in this thesis. Although extensive empirical research will be required to test the hypothetical bargaining hypothesis, the available data suggests that a version of hypothetical bargaining based on the benefit-equilibrating solution concept may offer an empirically relevant alternative explanation of how people identify mutually advantageous solutions of games with multiple Nash equilibria, and so the theory warrants further empirical testing.

It is important to note that although hypothetical bargaining is presented as a separate theory, the benefit-equilibrating solution concept suggested in this thesis is not incompatible with Sugden's (2011, 2015) version of the team reasoning theory based on the notion of mutual advantage: A 'team' can be viewed as a group of self-oriented decision-makers who act together in the attainment of an outcome which they identify as mutually advantageous by engaging in hypothetical bargaining. Note, however, that hypothetical bargaining is not compatible with those versions of the team reasoning theory which suggest that team reasoning requires group-identification based on overlapping interests and/or common experiences: Hypothetical bargaining can occur between complete strangers with incompatible personal interests.

On the surface of it, hypothetical bargaining is a parsimonious social coordination theory. It offers an explanation of how people coordinate their actions in a large variety of games with a relatively simple reasoning algorithm – a reasoning procedure which allows each decision-maker who uses it to identify a solution of a game from a commonly known information about decision-makers' preferences. It could also be viewed as providing a credible explanation of why one individual's deviation from the expected pattern of behaviour may trigger a negative response from other individuals: If hypothetical bargainers view a specific combination of players' strategies as an implicit agreement leading to the attainment of a mutually beneficial and agreeable outcome of a game, then each deviation from that strategy profile is viewed by them as a violation of their agreement.

According to the proponents of the theory, hypothetical bargaining is a rational social coordination theory. For this interpretation of hypothetical bargaining to be credible, a hypothetical bargainer's decision to play a part in the attainment of an outcome identified as the bargaining solution of a game must be shown to be compatible with the orthodox choice consistency principles. In this thesis, I have argued that a rational choice interpretation of hypothetical bargaining is conceptually problematic. The first reason

is hypothetical bargaining theory's vulnerability to choice rationalization problem: Common knowledge of the fact that each player is a hypothetical bargainer and has identified a certain outcome as the hypothetical bargaining solution of a game does not give the hypothetical bargainer a rational reason to play his or her part in the attainment of that outcome. Hypothetical bargainer's decision to play a part in the attainment of that outcome is only rational if s/he expects the other players to do that as well. The theory does not offer an explanation of why each hypothetical bargainer should hold such a belief, and so it does not provide an explanation of why s/he should be motivated to actually play his or her part in realizing an outcome which s/he recognizes as the hypothetical bargaining solution of a game.

I have argued that this difficulty could, in principle, be resolved if hypothetical bargaining were to be interpreted either as a choice or as a belief-formation algorithm, and suggested that lexicographic belief systems used in epistemic rationalizability models could, in principle, serve as an imperfect yet useful starting point for the development of an adequate formal representation of a belief-formation algorithm. However, both of the suggested interpretations of hypothetical bargaining involve non-trivial departures from the principles of orthodox game theory.

The second problem is that hypothetical bargaining, if interpreted as a rational-choice explanation of social coordination, can only account for the actions of individuals who express a common belief that the share of hypothetical bargainers in the population is sufficiently large. I have argued that such a belief can only form in a population where a sufficient number of individuals are acting as hypothetical bargainers. If interpreted as a rational-choice explanation of social coordination, the theory seems to fall into a vicious cycle: Hypothetical bargainers can only be motivated to act in a population state that cannot be reached without a sufficient number of players acting as hypothetical bargainers.

I have argued that the most plausible explanation of how such a population state could have emerged in the first place is that hypothetical bargaining first emerged as a familiarity heuristic which required little to no rational deliberation and was able to spread in the population due to its fitness-enhancing properties. However, evolutionary game theoretic models considered in this thesis only show that individuals using such a familiarity heuristic could be more successful than individuals using certain primitive decision rules. I have argued that the question of whether basic evolutionary game theory models are adequate for representing the actual evolution of decision-making rules cannot be answered with the tools of the evolutionary game theory itself. Therefore, evolutionary game theoretic explanation

suggested in this chapter is only valid insofar as it shows that the intuition that people's choices in non-cooperative games may be driven by an evolved sensitivity to benefit-allocation considerations has some conceptual credibility.

Finally, I have argued that hypothetical bargainers' ability to coordinate their actions may often depend on factors that are not related to payoff structures of games at all. Therefore, even if hypothetical bargaining could be developed into a complete and conceptually sound theory of strategic reasoning, a single generalizable model of hypothetical bargainers' *final* choices in non-cooperative games may not be possible after all.

The model of hypothetical bargaining based on the benefit-equilibrating solution concept which has been suggested and discussed in this thesis should not be viewed as a complete theory of social coordination, but rather as a conceptual exploration of the general principles of reasoning which may underly decision-makers' reasoning in various real-world interdependent decision problems, and as a study of how these principles could be incorporated into formal analysis of games. This thesis does not cover a considerable number of conceptually and empirically important questions that a complete theory of social coordination based on the notion of mutual advantage should be capable of answering.

Although the theoretical predictions of players' actions based on the benefit-equilibrating solution concept fits with some experimental findings from games discussed in this thesis, further empirical tests will need to be constructed to test the empirical validity of this account. One of the biggest empirical challenges is the problem of underdetermination. Further empirical research is needed to test competing theories about the modes of reasoning that people use in their interactions with each other. Such research will require sophisticated empirical tests specifically designed to distinguish the hypothetical bargaining theory from other competing approaches, such as team reasoning theory, cognitive hierarchy theory, social conventions theory, and coalitional rationalizability. Since individuals' actions can often be explained in terms of multiple accounts of what players try to achieve in games they play, these studies may need to consider a broader evidence base than mere observations of decision-makers' choices.

Another important question which has not been addressed in this thesis is what properties of decision-problems trigger individuals' search for mutually advantageous solutions of games. An adequate answer to this question will require an extensive empirical study of the structural and contextual features of decision problems in which people exhibit patterns of behaviour consistent with the theoretical predictions of the theory. Such empirical research could

answer some of the fundamental questions pertaining to the explanatory scope of the model, such as, for example, whether a behaviour consistent with the theoretical predictions of hypothetical bargaining can be observed in extensive form games.

One of the fundamental conceptual limitations of the hypothetical bargaining model discussed in this thesis is that it only provides a description of the formal properties that an outcome must have in order to be identified by the interacting decision-makers as the bargaining solution of a game. It does not offer a description of the process of mental bargaining that individuals go through in order to identify the bargaining solution. The actual details of the process of mental bargaining, such as players' perception of the decision problem and its context, as well as their beliefs about each other's perception of the decision problem, beliefs and motivations, may play a non-trivial role in the process of identification of hypothetical bargaining solutions. Therefore, there are serious reasons to believe that a highly idealized bargaining solutions may, at best, offer only approximately accurate predictions of hypothetical bargainers' actions. Further research of the psychological processes which underpin hypothetical bargaining may eventually lead to a development of a formal model of reasoning which could account for at least some of the factors which influence decision-makers' reasoning in real-world interdependent decision problems, and thus provide more accurate and testable predictions of hypothetical bargainers' actions. Further psychological research and, hopefully, a development of a reasonably accurate model of mental bargaining may ultimately be the only possible way to show that hypothetical bargaining is not merely a descriptive model which, like the other competing models, provides reasonably accurate predictions of people's choices, but as a model which provides an approximately accurate description of how people actually reason in non-cooperative games.

Bibliography

Alexander, J. M. [2007]: *The Structural Evolution of Morality*, Cambridge: Cambridge University Press.

Ambrus, A. [2009]: ‘Theories of Coalitional Rationality’, *Journal of Economic Theory*, 144(2), pp. 676-695.

Ambrus, A. [2006]: ‘Coalitional rationalizability’, *Quarterly Journal of Economics*, 121, pp. 903–929.

Ashcraft, M. H. [2006]: *Cognition*, Upper Saddle River, New Jersey: Pearson Education Inc.

Aumann, R. J., Brandenburger, A. [1995]: ‘Epistemic Conditions for Nash Equilibrium’, *Econometrica*, 63, pp. 1161-1180.

Aumann, R. J. [1987]: ‘Correlated Equilibrium as an Expression of Bayesian Rationality’, *Econometrica*, 55(1), pp. 1-18.

Aumann, R. J. [1985]: ‘What is Game Theory Trying to Accomplish?’, in K. Arrow and S. Honkapohja (eds.), ‘Frontiers of Economics’, Oxford: Basil Blackwell, pp. 28-76.

Aumann, R. J. [1976]: ‘Agreeing to Disagree’, *The Annals of Statistics*, 4(6), pp. 1236-1239.

Aumann, R. J. [1974]: ‘Subjectivity and Correlation in Randomized Strategies’, *Journal of Mathematical Economics*, 1, pp. 67-96.

Aumann, R. J. [1959]: ‘Acceptable Points in General Cooperative n -Person games’, in R. D. Luce and A. W. Tucker (eds.), ‘Contributions to the Theory of Games IV’, Princeton: Princeton University Press, pp. 287-324.

Bach, C., Perea, A. [2014]: ‘Utility Proportional Beliefs’, *International Journal of Game Theory*, 43, pp. 881-902.

- Bacharach, M. [2006]: *Beyond Individual Choice: Teams and Frames in Game Theory*, Princeton: Princeton University Press.
- Bacharach, M., Stahl D. O. [2000]: ‘Variable-Frame Level- n Theory’, *Games and Economic Behavior*, 32, pp. 220-246.
- Bacharach, M. [1999]: ‘Interactive Team Reasoning: A Contribution to the Theory of Co-operation’, *Research in Economics*, 53, pp. 117–147.
- Bacharach, M. and Bernasconi, M. [1997]: ‘The Variable Frame Theory of Focal Points: An Experimental Study’, *Games and Economic Behavior*, 19, pp. 1–45.
- Bacharach, M. [1993]: ‘Variable Universe Games’, in K. Binmore, A. Kirman and P. Tani (eds.), ‘Frontiers of Game Theory’, MIT Press.
- Bardsley, N., Mehta, J., Starmer, C. and Sugden, R. [2010]: ‘Explaining Focal Points: Cognitive Hierarchy Theory versus Team Reasoning’, *The Economic Journal*, 120, pp. 40–79.
- Bernheim, D. [1984]: ‘Rationalizable Strategic Behavior’, *Econometrica*, 52, pp. 1007-1028.
- Bicchieri, C. [2006]: *The Grammar of Society: The Nature and Dynamics of Social Norms*, New York: Cambridge University Press.
- Bicchieri [1995]: ‘The Epistemic Foundations of Nash Equilibrium’ in D. Little (ed.), ‘On the Reliability of Economic Models’, New York: Springer Science+Business Media, pp. 91-137.
- Binmore, K. [2009a]: *Rational Decisions*, Princeton: Princeton University Press.
- Binmore, K. [2009b]: ‘Interpersonal Comparison of Utility’, in D. Ross and H. Kincaid (eds.), ‘The Oxford Handbook of Philosophy of Economics’, Oxford University Press, pp. 540–559.
- Binmore, K. [2008]: ‘Do Conventions Need to Be Common Knowledge?’, *Topoi* 27(1), pp. 17-27.
- Binmore, K. [2005]: *Natural Justice*, Oxford and New York: Oxford University Press.
- Binmore, K. [1992]: *Fun and Games: A Text on Game Theory*, D. C. Heath and Company, Lexington, MA.

- Binmore, K., Rubinstein, A., Wolinsky, A. [1986]: 'The Nash Bargaining Solution in Economic Modelling', *RAND Journal of Economics*, pp. 176-188.
- Binmore, K. [1980]: 'Nash Bargaining Theory I', International Center for Economics and Related Disciplines, London School of Economics, Discussion Paper 80-09.
- Blume, L. E., Brandenburger, A., Dekel, E., [1991a]: 'Lexicographic Probabilities and Choice under Uncertainty', *Econometrica* 59, pp. 61-79.
- Blume, L. E., Brandenburger, A., Dekel, E., [1991b]: 'Lexicographic Probabilities and Equilibrium Refinements', *Econometrica* 59, pp. 81-98.
- Brandenburger, A., Dekel, E. [1987]: 'Rationalizability and Correlated Equilibria', *Econometrica* 55(6), pp. 1391-1402.
- Cabon-Dhersin, M., Etchart-Vincent N. [2012]: 'The Puzzle of Cooperation in a Game of Chicken: An Experimental Study', *Theory and Decision*, 72, pp. 65-87.
- Camerer, C. F., Ho, T., Chong, J. [2004]: 'A Cognitive Hierarchy Model of Games', *Quarterly Journal of Economics*, 119(3), pp. 861-898.
- Chaudhuri, A. [2011]: 'Sustaining Cooperation in Laboratory Public Goods Experiments: A Selective Survey of the Literature', *Experimental Economics*, 14, pp. 47-83.
- Chen, X., Deng, X, Teng, S. [2009]: 'Settling the Complexity of Computing Two-Player Nash Equilibria', *Journal of the ACM*, 56(3).
- Colman, A. M., Pulford, B. D., and Lawrence, C. L. [2014]: 'Explaining Strategic Coordination: Cognitive Hierarchy Theory, Strong Stackelberg Reasoning, and Team Reasoning', *Decision*, 1, pp. 35-58.
- Colman, A. M., Pulford, B. D., and Rose, J. [2008]: 'Collective Rationality in Interactive Decisions: Evidence for Team Reasoning', *Acta Psychologica*, 128, pp. 387-397.
- Colman, A. M. and Stirk, J. A. [1998]: 'Stackelberg Reasoning in Mixed-Motive Games: An Experimental Investigation', *Journal of Economic Psychology*, 19, pp. 279-293.

- Colman, A. M. [1997]: ‘Salience and Focusing in Pure Coordination Games’, *Journal of Economic Methodology*, 4(1), pp. 61-81.
- Colman, A. M., Bacharach, M. [1997]: ‘Payoff Dominance and the Stackelberg Heuristic’, *Theory and Decision*, 43, pp. 1-19.
- Conley, J. P., Wilkie, S. [2012]: ‘The Ordinal Egalitarian Bargaining Solution for Finite Choice Sets’, *Social Choice and Welfare*, 38, pp. 23-42.
- Costa-Gomes, M. A., Crawford, V. P., Broseta, B. [2001]: ‘Cognition and Behavior in Normal-Form Games: An Experimental Study’, *Econometrica*, 69(5), pp. 1193–1235.
- Crawford, V. P., Gneezy, U., Rottenstreich, Y. [2008]: ‘The Power of Focal Points is Limited: Even Minute Payoff Asymmetry May Yield Large Coordination Failures’, *The American Economic Review*, 98, pp. 1443-1458.
- Cressman, R. [1997]: ‘Local Stability of Smooth Selection Dynamics for Normal Form Games’, *Mathematical Social Sciences*, 34, pp. 1–19.
- Cubitt, R., Sugden, R. [2003]: ‘Common Knowledge, Salience, and Convention: A Reconstruction of David Lewis’s Game Theory’, *Economics and Philosophy*, 19, pp. 175–210.
- van Damme, E. [1987]: *Stability and Perfection of Nash Equilibria*, Berlin: Springer-Verlag.
- D’Arms, J. [1996]: ‘Sex, Fairness, and the Theory of Games’, *Journal of Philosophy*, 93 (12), pp. 615–627.
- Dekel, E., Siniscalchi, M. [2015]: ‘Epistemic Game Theory’, in H. Peyton-Young and S. Zamir (eds.), ‘Handbook of Game Theory with Economic Applications, Vol. 4’, Amsterdam and Oxford: Elsevier, pp. 619-702.
- Dietrich, F., List, C. [2016]: ‘Mentalism Versus Behaviourism in Economics: A Philosophy-of-Science Perspective’, *Economics and Philosophy*, 32, pp. 249–281.
- Ellsberg, D. [1956]: ‘A Theory of the Reluctant Duelist’, *American Economic Review*, 46, pp. 909-923.

- Elster, J. [1989]: *The Cement of Society: A Study of Social Order*, Cambridge: Cambridge University Press.
- Faillo, M., Smerilli, A., and Sugden, R. [2016]: ‘Can a Single Theory Explain Coordination? An Experiment on Alternative Modes of Reasoning and the Conditions Under Which They Are Used’, CBES [Centre for Behavioural and Experimental Social Science] Working paper 16–01, University of East Anglia.
- Faillo, M., Smerilli, A. and Sugden, R. [2013]: ‘The Roles of Level-k and Team Reasoning in Solving Coordination Games’, Cognitive and Experimental Economics Laboratory Working Paper (No. 6-13), Department of Economics, University of Trento, Italy.
- Fehr, E. and Schmidt, K. M. [1999]: ‘A Theory of Fairness, Competition, and Cooperation’, *Quarterly Journal of Economics*, 114, pp. 817-868.
- Fudenberg, D., Levine, D. K., [1998]: *The Theory of Learning in Games*, Cambridge, MA: The MIT Press.
- Fudenberg, D. and Tirole, J. [1991]: *Game Theory*, Massachusetts: The MIT Press.
- Gauthier, D. [2013]: ‘Twenty-Five On’, *Ethics*, 123, pp. 601–624.
- Gauthier, D. [1986]: *Morals By Agreement*, Oxford: Oxford University Press.
- Gauthier, D. [1975]: ‘Coordination’, *Dialogue*, 14, pp. 195-221.
- Gigerenzer, G., Selten, R. (eds.) [2001]: *Bounded Rationality: The Adaptive Toolbox*. Cambridge, MA: The MIT Press.
- Gigerenzer, G., Todd, P. M., the ABC Research Group [1999]: *Simple Heuristics That Make Us Smart*. Oxford: Oxford University Press.
- Gilbert, M. [1989]: ‘Rationality and Salience’, *Philosophical Studies*, 57, pp. 61-77.
- Gilboa, I., Matsui, A. [1991]: ‘Social Stability and Equilibrium’, *Econometrica*, 59, pp. 859–867.
- Gintis, H. [2008]: *The Bounds of Reason: Game Theory and the Unification of the Behavioral Sciences*, Princeton and Oxford: Princeton University Press.

- Gold, N. [2012]: ‘Team Reasoning, Framing and Cooperation’, in S. Okasha and K. Binmore (eds.), ‘Evolution and Rationality: Decisions, Co-Operation and Strategic Behaviour’, Cambridge University Press, Cambridge, pp. 185–212.
- Gold, N., Sugden, R. [2007a]: ‘Collective Intentions and Team Agency’, *Journal of Philosophy*, 104, pp. 109–137.
- Gold, N., Sugden, R. [2007b]: ‘Theories of Team Agency’, in F. Peter and H. B. Schmidt (eds.), ‘Rationality and Commitment’, Oxford: Oxford University Press, pp. 280–312.
- Gold, N., List, C. [2004]: ‘Framing as Path-Dependence’, *Economics and Philosophy*, 20, pp. 253–277.
- Guala, F. [2013]: ‘The Normativity of Lewis Conventions’, *Synthese*, pp. 3107–3122.
- Hammond, P. J. [1991]: ‘Interpersonal Comparisons of Utility: Why and How They Are and Should Be Made’, in J. Elster and J. E. Roemer (eds.), ‘Interpersonal Comparisons of Well-Being’, Cambridge: Cambridge University Press, pp. 200–254.
- Harsanyi J., Selten R. [1988]: *A General Theory of Equilibrium Selection in Games*, Cambridge, MA: The MIT Press.
- Harsanyi, J. [1968a]: ‘Games with Incomplete Information Played by “Bayesian” Players. Part II: Bayesian Equilibrium Points’, *Management Science*, 14, pp. 320–334.
- Harsanyi, J. [1968b]: ‘Games with Incomplete Information Played by “Bayesian” Players. Part III: The Basic probability Distribution of the Game’, *Management Science*, 14, pp. 486–502.
- Harsanyi, J. [1967]: ‘Games with Incomplete Information Played by “Bayesian” Players. Part I: The Basic Model’, *Management Science*, 14, pp. 159–182.
- Hausman, D. M. [2012]: *Preference, Value, Choice, and Welfare*, Cambridge University Press, New York, NY.
- Hausman, D. M. (1995): ‘The Impossibility of Interpersonal Utility Comparisons’, *Mind*, 104, pp. 473–490.

- Hines, W. G. S. [1980]: ‘Three characterizations of population strategy stability’, *Journal of Applied Probability*, 17, pp. 333–340.
- Ho, T., Camerer, C. and Weigelt, K. [1998]: ‘Iterated Dominance and Iterated Best Response in Experimental ‘p-Beauty Contests’, *American Economic Review*, 88(4), pp. 947–969.
- Hofbauer, J., Sigmund, K. [1998]: *Evolutionary Games and Population Dynamics*. Cambridge: Cambridge University Press.
- Hofbauer, J. [1995a]: ‘Imitation Dynamics for Games’, University of Vienna, [unpublished paper].
- Hofbauer, J. [1995b], ‘Stability for the Best Response Dynamics’, University of Vienna [unpublished paper].
- Hopkins, E. [1999]: ‘A Note on Best-Response Dynamics’, *Games and Economic Behavior*, 29, pp. 138-150.
- Hurley, S. [2005a]: ‘Rational Agency, Cooperation and Mind-Reading’, in N. Gold (ed), ‘Teamwork: Multi-Disciplinary Perspectives’, New York: Palgrave MacMillan.
- Hurley, S. [2005b]: ‘Social Heuristics that Make Us Smarter’, *Philosophical Psychology*, 18, pp. 585-612.
- Infante, G., Lecouteux G., Sugden R. [2016]: Preference Purification and the Inner Rational Agent: A Critique of the Conventional Wisdom of Behavioural Welfare Economics. *Journal of Economic Methodology*, 23 (1), pp. 1-25.
- Janssen, M. C. W. [2001]: ‘Rationalizing Focal Points’, *Theory and Decision*, 50, pp. 119-148.
- Kajii, A., Morris, S. [1997]: ‘Common p -Belief: The General Case’, *Games and Economic Behaviour*, 18, pp. 73-82.
- Kalai, E. [1977]: ‘Proportional Solutions to Bargaining Situations: Interpersonal Utility Comparisons’, *Econometrica*, 45, pp. 1623–1630.
- Kalai, E. and Smorodinsky, M. [1975]: ‘Other Solutions to Nash’s Bargaining Problem’, *Econometrica*, 43, pp. 513–518.
- Karpus, J. and Radzvilas, M. [2016]: ‘Team Reasoning and a Measure of Mutual Advantage in Games’ [accepted to *Economics and Philosophy*].

- Kets, W., Sandroni A. [2014]: ‘A Belief-Based Theory of Homophily’ [working paper].
- Kreps, D. M. [1990]: *A Course in Microeconomic Theory*, Harlow: Prentice Hall.
- Kreps, D. [1988]: *Notes on the Theory of Choice*, Boulder: Westview Press, Inc.
- Ledyard, J. O. [1995]: ‘Public Goods: A Survey of Experimental Research’, in J. H. Kagel and A. E. Roth (eds.), ‘The Handbook of Experimental Economics’, Princeton University Press, Princeton, pp. 111–194.
- Lewis, D. [1969]: *Convention: A Philosophical Study*, Cambridge: Harvard University Press.
- Little, I. M. D. [1949]: ‘A Reformulation of the Theory of Consumers’ Behaviour’, *Oxford Economic Papers*, 1, pp. 90–99.
- Luce, R. D. and Raiffa, H. [1957]: *Games and Decisions: Introduction and Critical Survey*, New York: Dover Publications, Inc.
- Luo X., Yang, C. [2009]: ‘Bayesian Coalitional Rationalizability’, *Journal of Economic Theory*, 144(1), pp. 248–263.
- Maynard Smith, J., Price, G. R. [1973]: ‘The logic of animal conflict’. *Nature*, 246 (5427), pp. 15–18.
- Maschler, M., Solan, E., Zamir, S. [2013]: *Game Theory*, Cambridge: Cambridge University Press.
- Matsui, A. [1992]: ‘Best-Response Dynamics and Socially Stable Strategies’, *Journal of Economic Theory*, pp. 343–362.
- Mech, D. L. [1999]: ‘Alpha Status, Dominance, and Division of Labor in Wolf Packs’, *Canadian Journal of Zoology*, 77, pp. 1196–1203.
- Mehta, J., Starmer, C., Sugden R. [1994]: ‘Focal Points in Pure Coordination Games: An Experimental Investigation’, *Theory and Decision*, 36(2), pp. 163–185.
- Misyak, J., Chater, N. [2014]: ‘Virtual Bargaining: A Theory of Social Decision-Making’, *Philosophical Transactions of the Royal Society, B* 369, pp. 1–9.

- Misyak, J. B., Melkonyan, T., Zeitoun, H., Chater, N. [2014]: ‘Unwritten Rules: Virtual Bargaining Underpins Social Interaction, Culture, and Society’, *Trends in Cognitive Sciences*, 18, pp. 512-519.
- Myerson, R. B. [1991]: *Game Theory: Analysis of Conflict*, Cambridge and London: Harvard University Press.
- Myerson, R. B. [1977]: ‘Two-Person Bargaining Problems and Comparable Utility’, *Econometrica*, 45, pp. 1631–1637.
- Monderer, D., Samet, D. [1989]: ‘Approximating Common Knowledge with Common Beliefs’, *Games and Economic Behaviour*, pp. 170-190.
- Moreno, D., Wooders, J. [1996]: ‘Coalition-Proof Equilibrium’, *Games and Economic Behavior*, 17, pp. 80-112.
- Morris, S. [1995]: ‘The Common Prior Assumption in Economic Theory’, *Economics and Philosophy*, 11, pp. 227-253.
- Nagel, R., [1995]: ‘Unraveling in Guessing Games: An Experimental Study’, *American Economic Review*, 85(5): pp. 1313–1326.
- Nash, J. [1951]: ‘Non-Cooperative Games’, *The Annals of Mathematics*, 54(2), pp. 286-295.
- Nash, J. [1950a]: ‘The Bargaining Problem’, *Econometrica*, 18(2), pp. 155-162.
- Nash, J. [1950b]: ‘Equilibrium points in n -person games’, *Proceedings of the National Academy of Sciences*, 36(1), pp. 48-49.
- Nisan, N., Schapira, M., Valiant, G., Zohar, A. [2011]: *Best-Response Mechanisms*, ICS-11: Proceedings of the 2nd Symposium Conference on Innovations in Computer Science (ICS), pp. 155-165.
- Nydegger, R. and Owen, G. [1974]: ‘Two-Person Bargaining: An Experimental Test of the Nash Axioms’, *International Journal of Game Theory*, 3, pp. 239–249.
- Olcina, G. and Urbano, A. [1994]: ‘Introspection and Equilibrium Selection in 2x2 Matrix Games’, *International Journal of Game Theory*, 23, pp. 183-206.
- Osborne, M., J., Rubinstein, A. [1994]: *A Course in Game Theory*, Cambridge and London: The MIT Press.

- Pearce, D. [1984]: ‘Rationalizable Strategic Behavior and the Problem of Perfection’, *Econometrica*, 52, pp. 1029-1050.
- Pearl, Judea [1983]: *Heuristics: Intelligent Search Strategies for Computer Problem Solving*, New York, Addison-Wesley.
- Perea, A. [2014]: ‘From Classical to Epistemic Game Theory’, *International Game Theory Review*, 16(1), pp. 1-22.
- Perea, A. [2012]: *Epistemic Game Theory: Reasoning and Choice*, Cambridge: Cambridge University Press.
- Perea, A. [2011]: ‘An Algorithm for Proper Rationalizability’, *Games and Economic Behavior*, 72, pp. 510-525.
- Raiffa, H. [1953]: ‘Arbitration Schemes for Generalized Two Person Games’, in H. W. Kuhn, and A. W. Tucker (eds.), ‘Contributions to the Theory of Games II’, Princeton: Princeton University Press, pp. 361-387.
- Roth, A. [1979]: ‘Proportional Solutions to the Bargaining Problem’, *Econometrica*, 47, pp. 775-778.
- Roth, A. [1977]: ‘Individual Rationality and Nash’s Solution to the Bargaining Problem’, *Mathematics of Operations Research*, 2(1), pp. 64-65.
- Rubinstein, A. and Salant, Y. [2016]: ‘“Isn’t Everyone Like Me?”: On the Presence of Self-Similarity in Strategic Interactions’, *Judgment and Decision Making*, 11, pp. 168–173.
- Rubinstein, A. and Salant, Y. [2008]: ‘Some Thoughts on the Principle of Revealed Preference’, in A. Caplin and A. Schotter, (eds.), ‘The Foundations of Positive and Normative Economics: A Handbook’, New York: Oxford University Press, pp. 116–124.
- Rubinstein, A. [1991]: ‘Comments on the Interpretation of Game Theory’, *Econometrica*, 59, pp. 909-924.
- Rubinstein, A. [1982]: ‘Perfect Equilibrium in a Bargaining Model’, *Econometrica*, 50, pp. 97-109.
- Samet, D. [2013]: ‘Common Belief in Rationality in Games of Perfect Information’, *Games and Economic Behavior*, 79, pp. 192-200.

- Samuelson, P. [1938]: ‘A Note on the Pure Theory of Consumers’ Behaviour’, *Economica* 5, pp. 61–71.
- Sandholm, W. H., Dokumaci, E., Franchetti, F. [2012]: ‘Dynamo: Diagrams for Evolutionary Games Dynamics’, <http://www.ssc.wisc.edu/~whs/dynamo>.
- Sandholm, W. H. [2010a]: *Population Games and Evolutionary Dynamics*, Cambridge, MA and London: The MIT Press.
- Sandholm, W. H. [2010b]: ‘Local Stability Under Evolutionary Game Dynamics’, *Theoretical Economics*, 5, pp. 27–50.
- Savage, L. J. [1954]: *The Foundation of Statistics*, New York: Wiley.
- Schelling, T. [1960]: *The Strategy of Conflict*, Cambridge: Harvard University Press.
- Selten, R. [1988]: *Models of Strategic Rationality*. Dordrecht: Springer Science+Business Media Dordrecht.
- Skyrms, B. [1996]: *Evolution of the Social Contract*, Cambridge: Cambridge University Press.
- Skyrms, B. [1990]: ‘The Dynamics of Rational Deliberation’, Cambridge, MA and London: The MIT Press.
- Smerilli, A. [2014]: ‘Theories of Team Reasoning’, *Philosophical Readings*, VI(3), pp. 24–34.
- Smerilli, A. [2012]: ‘We-Thinking and Vacillation Between Frames: Filling a Gap in Bacharach’s Theory’, *Theory and Decision*, 73, pp. 539–560.
- Stahl, D. O., II, and Paul W. Wilson, P. W. [1995]: ‘On Players’ Models of Other Players: Theory and Experimental Evidence’, *Games and Economic Behavior*, 10(1), pp. 218–54.
- Stahl, D. O., II, Wilson, P. W. [1994]: ‘Experimental Evidence on Players’ Models of Other Players’, *Journal of Economic Behavior and Organization*, 25(3), pp. 309–27.
- Sugden, R. [2015]: ‘Team Reasoning and Intentional Cooperation for Mutual Benefit’, *Journal of Social Ontology*, 1, pp. 143–166.
- Sugden, R. [2011]: ‘Mutual Advantage, Conventions and Team Reasoning’, *International Review of Economics*, 58, pp. 9–20.

- Sugden, R. [2003]: ‘The Logic of Team Reasoning’, *Philosophical Explorations: An International Journal for the Philosophy of Mind and Action*, 6, pp. 165–181.
- Sugden, R. [2000]: ‘Team Preferences’, *Economics and Philosophy*, 16, pp. 175–204.
- Sugden, R. [1993]: ‘Thinking as a Team: Towards an Explanation of Nonselfish Behavior’, *Social Philosophy and Policy*, 10, pp. 69–89.
- Sugden, R. [1995]: ‘A Theory of Focal Points’, *Economic Journal*, 105, pp. 533–550.
- Sugden, R. [1986]: *The Economics of Rights, Co-operation, and Welfare*, New York: Palgrave Macmillan.
- Taylor, P. D., Jonker, L. [1978]: Evolutionarily Stable Strategies and Game Dynamics. *Mathematical Biosciences* 40, pp. 145–156.
- Thomas, B. [1984]: ‘Evolutionary Stability: States and Strategies’, *Theoretical Population Biology* 26 (1), pp. 49–67.
- Tversky, A., Kahneman, D. [1974]: ‘Judgment under Uncertainty: Heuristics and Biases’, *Science*, 185 (4157), pp. 1124–1131.
- Vanderschraaf, P. [2001]: *Learning and Coordination: Inductive Deliberation, Equilibrium, and Convention*, London: Routledge.
- Vanderschraaf, P. [1998]: ‘Knowledge, Equilibrium, and Convention’, *Erkenntnis*, 49, pp. 337–369.
- Vanderschraaf, P. [1995]: ‘Convention as Correlated Equilibrium’, *Erkenntnis*, 42, pp. 65–87.
- Weibull, J. [1995]: *Evolutionary Game Theory*, Cambridge, MA and London: The MIT Press.
- Zizzo, D. J. and Tan, J. H. [2007]: ‘Perceived Harmony, Similarity and Cooperation in 2X2 Games: An Experimental Study’, *Journal of Economic Psychology*, 28, pp. 365–386.
- Zhang, D., Zhang, Y. [2008]: ‘An Ordinal Bargaining Solution with Fixed-Point Property’, *Journal of Artificial Intelligence Research*, 33, pp. 433–464.